ADVANCED TECHNIQUES (MC/MD)

A (seemingly) random selection.

Daan Frenkel

Beyond Newtonian MD

1. Langevin dynamics
2. Brownian dynamics
3. Stokesian dynamics
4. Dissipative particle dynamics
5. Etc. etc.

WHY?

1. Can be used to simulate molecular motion in a viscous medium, *without solving the equations of motion for the solvent particles.*
2. Can be used as a *thermostat*.

Friction force

Conservative force

First, consider m ion with frictio alone:

$$\xcancel{m\dot{\vec{v}} = -\gamma\vec{v}(t') - \nabla U \ ,}$$

After a short while, all particles will stop moving, due to the fricti

“random” force

Better:

$$m\dot{\vec{v}} = -\gamma\vec{v}(t') - \nabla U + \vec{\zeta}(t) \ ,$$

There is a relation between the correlation function of the random force and the friction coefficient:

$$< \zeta_x(0)\zeta_x(t) >= 2kT\gamma\delta(t)$$

The derivation is straightforward, but beyond the scope of this lecture.

The KEY point is that the friction force and the random force ARE RELATED.

Limiting case of Langevin dynamics:

No inertial effects (m=0)

$$m\dot{\vec{v}} = -\gamma\vec{v}(t') - \nabla U + \vec{\zeta}(t) \ ,$$

Becomes:

$$0 = -\gamma\vec{v}(t') - \nabla U + \vec{\zeta}(t) \ ,$$

“Brownian Dynamics”

(But still the friction force and the random force are related)

What is missing in Langevin dynamics and Brownian dynamics?

1. Momentum conservation
2. Hydrodynamics

(1 and 2 are not independent).

Is this serious?

Not always: it depends on the time scales.

Momentum “diffuses” away in a time L^2/ν. After that time, a “Brownian” picture is OK.

However: hydrodynamics makes that the friction constant depends on the positions of all particles (and so do the random forces…).

Momentum conserving, coarse-grained schemes:

- Dissipative particle dynamics
- Stochastic Rotation Dynamics

These schemes represent the solvent explicitly (i.e. as particles), but in a highly simplified way.

ADVANCED MC SAMPLING

Is the rejection of Monte Carlo trial moves wasteful?

Conventional MC performs a random walk in configuration space, such that the number of times that each point is visited, is proportional to its Boltzmann weight.

$$n(\mathbf{r}^N) = c\,\exp[-\beta\mathcal{U}(\mathbf{r}^N)]$$

$$\frac{\text{acc}(o \to n)}{\text{acc}(n \to o)} = \frac{\mathcal{N}(n)}{\mathcal{N}(o)} = \exp\{-\beta[\mathcal{U}(n) - \mathcal{U}(o)]\}$$

Metropolis,
Rosenbluth,Rosenbluth,

Teller and Teller choice:

$$\text{acc}(o \to n) = \min\left(1, \exp\{-\beta[\mathcal{U}(\mathbf{r}'^N) - \mathcal{U}(\mathbf{r}^N)]\}\right)$$

Solution of conflict: play with the a-priori probabilities of trial moves:

$$\alpha(o \rightarrow n) \neq \alpha(n \rightarrow o)$$

$$\frac{\mathrm{acc}(o \to n)}{\mathrm{acc}(n \to o)} = \frac{\alpha(n \to o)}{\alpha(o \to n)} \exp\{-\beta[\mathcal{U}(n) - \mathcal{U}(o)]\}.$$

In particular, if:

$$\frac{\alpha(n \to o)}{\alpha(o \to n)} = \exp\{-\beta[\mathcal{U}(o) - \mathcal{U}(n)]\}.$$

Then

$$\frac{\mathrm{acc}(o \to n)}{\mathrm{acc}(n \to o)} = 1 \qquad \text{(100\% acceptance)}$$

100% acceptance can be achieved in special cases: e.g. Swendsen-Wang, Wolff, Luyten, Whitelam-Geissler, Bortz-Kalos-Lebowitz, Krauth…

General idea: construct “cluster moves”

Simplest example: Swendsen-Wang

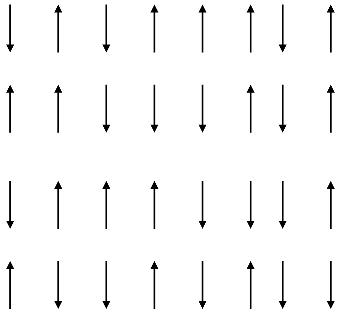

Illustration: 2D Ising model.

Snapshot: some neighbors are parallel, others anti-parallel

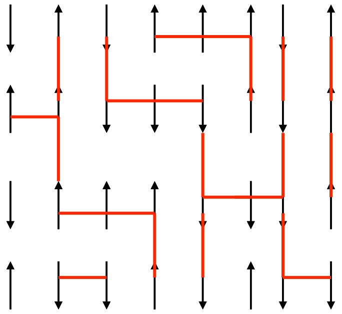

Number of parallel nearest-neighbor pairs: $\mathbf{N_p}$

Number of anti-parallel nearest neighbor pairs is: $\mathbf{N_a}$

Total energy: $\mathbf{U = (N_a - N_p)\,J}$

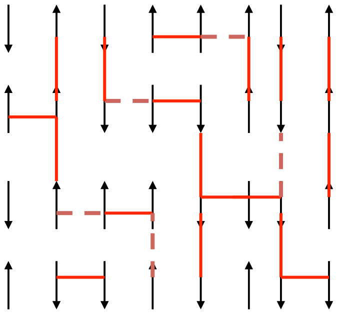

Make “bonds” between parallel neighbors. The probability to have a bond (red line) between parallel neighbors is **p** (as yet undetermined). With a probability **1-p**, parallel neighbors are not connected (blue dashed line).

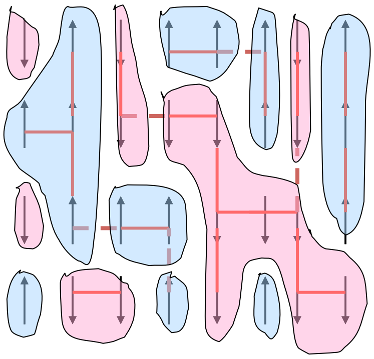

Form clusters of all spins that are connected by bonds. Some clusters are all “spin up” others are all “spin down”.

Denote the number of clusters by **M**.

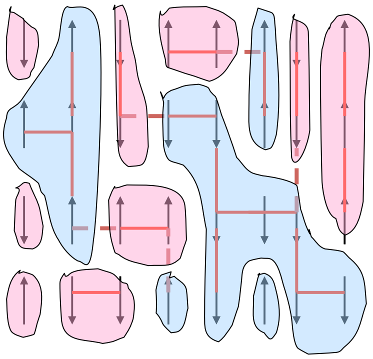

Now randomly flip clusters. This yields a new cluster configuration with probability **$P_{(flip)} = (1/2)^M$.**

Then reconnect parallel spins

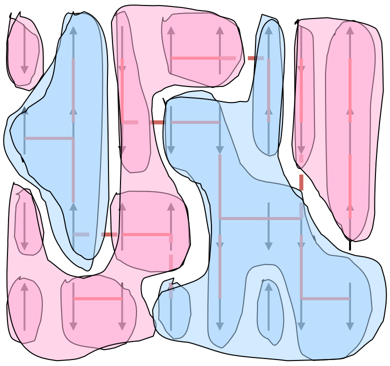

Next: forget about the “bonds”…

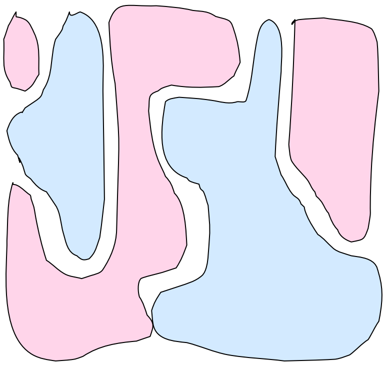

New spin configuration!

$$P_o P_{clus}(o) P_{flip}(M) P_{acc}(o \rightarrow n)$$

$$=$$

$$P_n P_{clus}(n) P_{flip}(M) P_{acc}(n \rightarrow o)$$

$$\exp(-\beta U_o) p^{n_c} (1-p)^{N_p(o)-n_c} (1/2)^M P_{acc}(o \rightarrow n)$$

$$=$$

$$\exp(-\beta U_n) p^{n_c} (1-p)^{N_p(n)-n_c} (1/2)^M P_{acc}(n \rightarrow o)$$

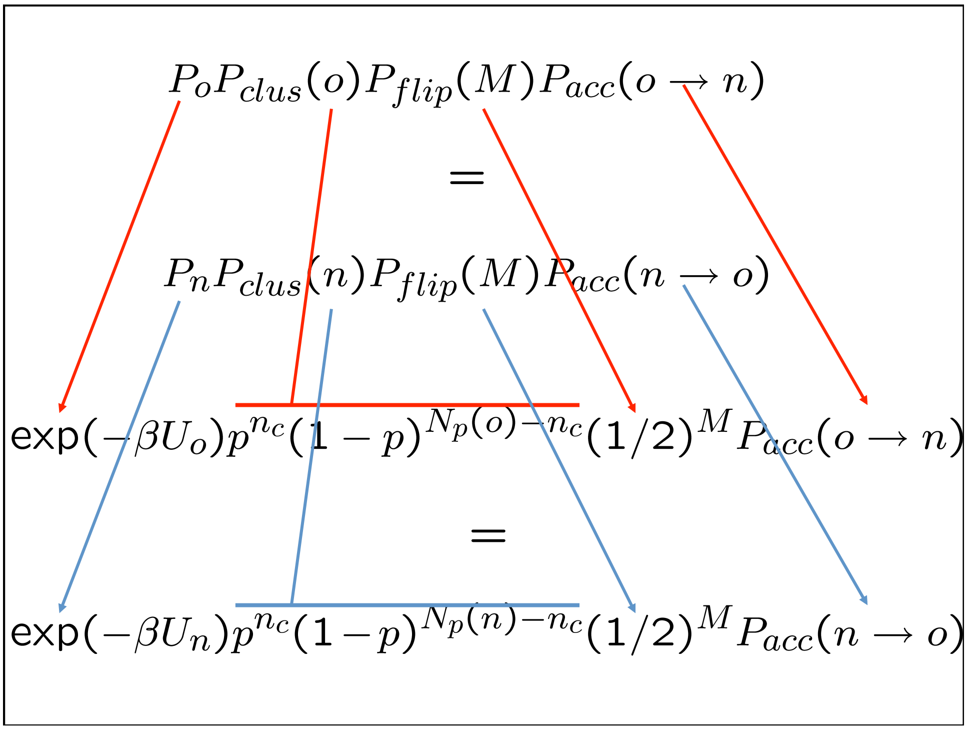

$$P_o P_{clus}(o) P_{flip}(M) P_{acc}(o \to n)$$

$$=$$

$$P_n P_{clus}(n) P_{flip}(M) P_{acc}(n \to o)$$

$$\exp(-\beta U_o) p^{n_c} (1-p)^{N_p(o)-n_c} (1/2)^M P_{acc}(o \to n)$$

$$=$$

$$\exp(-\beta U_n) p^{n_c} (1-p)^{N_p(n)-n_c} (1/2)^M P_{acc}(n \to o)$$

Moreover, we want 100% acceptance, i.e.:

$P_{acc}(o \to n) = P_{acc}(n \to o) = 1$

$$\exp(-\beta U_o) \cancel{p^{n_c}} (1-p)^{N_p(o)-\cancel{n_c}} \cancel{(1/2)^M} P_{acc}(o \to n)$$

$$=$$

$$\exp(-\beta U_n) \cancel{p^{n_c}} (1-p)^{N_p(n)-\cancel{n_c}} \cancel{(1/2)^M} P_{acc}(n \to o)$$

Hence:

$$\exp(-\beta U_o)(1-p)^{N_p(o)} = \exp(-\beta U_n)(1-p)^{N_p(n)}$$

$$\exp(\beta(U_n - U_o)) = (1-p)^{N_p(n)-N_p(o)}$$

But remember:

$$U_n - U_o = J(N_a(n) - N_p(n)) - J(N_a(o) - N_p(o))$$

or

$$\Delta U = J(\Delta N_a - \Delta N_p)$$

But: $\Delta N_a = -\Delta N_p$

and therefore

$$\Delta U = -2J\Delta N_p$$

$$\exp(\beta(U_n - U_o)) = \exp(-2\beta J(N_p(n) - N_p(o)))$$

Combining this with:

$$\exp(\beta(U_n - U_o)) = (1-p)^{N_p(n)-N_p(o)}$$

we obtain:

$$p = 1 - \exp(-2\beta J)$$

100% acceptance!!!

WASTE RECYCLING MC

Include “rejected” moves in the sampling

This is the key:

$$\sum_m \rho(m)\pi_{mn} = \rho(n)$$

The transition matrix π leaves the equilibrium distribution ρ unchanged.

$$\langle A \rangle_\rho = \sum_n A_n \rho_n$$

This, we can rewrite as:

$$\sum_n A_n \rho_n = \sum_n \sum_m A_n \rho_m \pi_{mn} = \sum_m \rho_m \sum_n A_n \pi_{mn}$$

$$= \sum_m \rho_m \sum_n A_n \pi_{mn} \Leftrightarrow \langle A \rangle_\rho = \left\langle \sum_n \pi_{mn} A_n \right\rangle_{\rho_m}$$

$$\langle A \rangle_{\rho} = \left\langle \sum_n \pi_{mn} A_n \right\rangle_{\rho_m}$$

Note that <A> is no longer an average over “visited” states – we also include “rejected” moves in the sampling.

Slightly dishonest and slightly trivial example:

Sampling the magnetization of a 2D Ising system

Compare:

1. Normal (Swendsen-Wang) MC (sample **one** out of 2^n states)
2. Idem + “waste recycling” (sample **all** 2^n states)

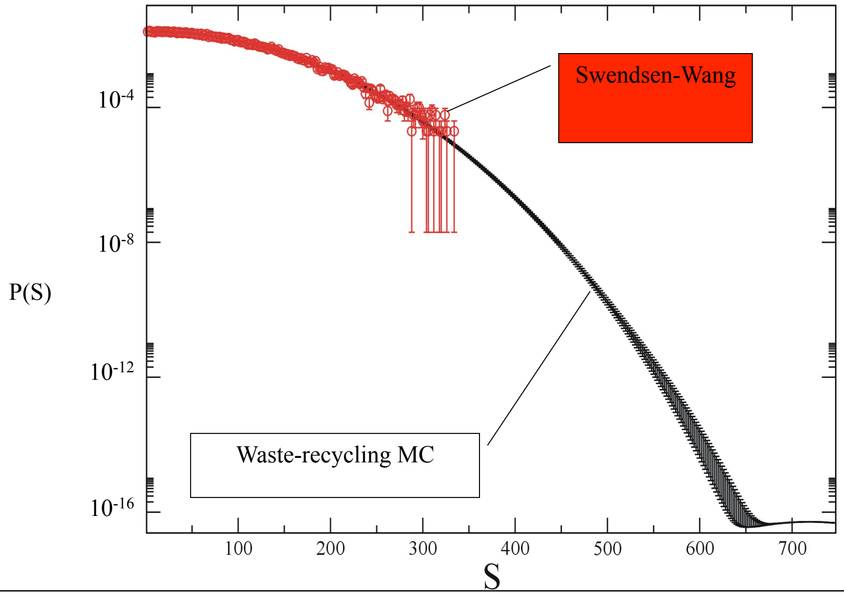

Monte Carlo sampling with noisy weight functions.

Two possible cases:

1. **The calculation of the energy function is subject to statistical error** (Ceperley, Dewing, J. Chem. Phys. **110**, 9812 (1999).)

$$u_{\text{computed}} = u_{\text{real}} + \delta u$$

with:

$$< \delta u >= 0$$
$$< (\delta u)^2 >= \sigma_s^2$$

We will assume that the fluctuations in ***u*** are Gaussian. Then:

Now consider that we do Monte Carlo with this noisy energy function:

$$\frac{P_n(\mathbf{x}_n)}{P_o(\mathbf{x}_o)} = \exp[-\beta \Delta u]$$

with

$$\Delta u = u_n + \delta u_n - u_o - \delta u_o$$

Then:

$$\left\langle \frac{P_n}{P_o} \right\rangle = \exp[-\beta \langle \Delta u \rangle + (\beta\sigma)^2/2]$$

With: $\sigma^2 = 2\sigma_s^2$

As a consequence, we sample the states with the wrong weight.

However, we can use another acceptance rule:

$$P_{\text{acc}} = \text{Min}\{1, \exp[-\beta\Delta u - (\beta\sigma)^2/2]\}$$

In that case:

$$\left\langle \frac{P_n}{P_o} \right\rangle = \exp[-\beta\langle\Delta u\rangle + (\beta\sigma)^2/2] \times \exp[-(\beta\sigma)^2/2]$$

$$= \exp[-\beta\langle\Delta u\rangle]$$

In other words:

If the statistical noise in the energy is Gaussian,

and its variance is constant,

then we can perform rigorous sampling, even when the energy function is noisy

2. **The weight function is noisy, but its average is correct** (not so common in molecular simulation, but quite common in other sampling problems)

(can also be sampled rigorously – but outside the scope of this lecture)

Recursive sampling

Outline:

1. Recursive enumeration
    a) Polymer statistics (simulation)
    b) ..
2. Molecular Motors (experiments!)
    (well, actually, simulated experiments)

Lattice polymers:

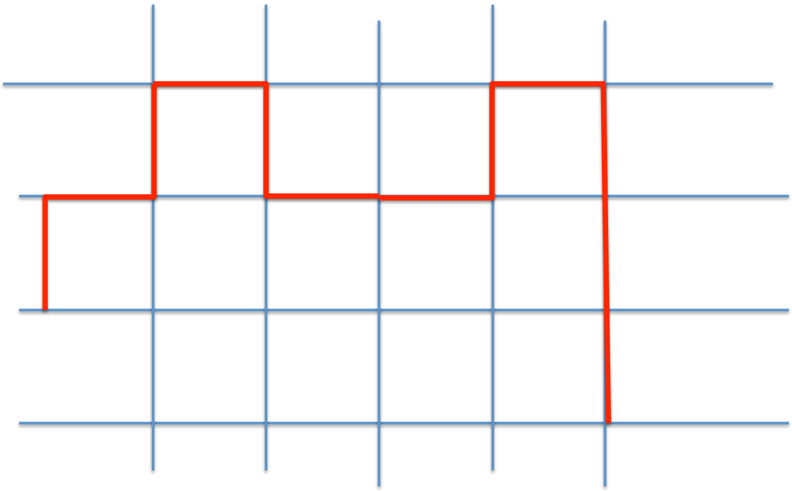

Consider a lattice (e.g. 2D-square).

At a given point x_i, the potential energy is $U(x_i)$.

The Boltzmann factor for a particle at point x_i is

$\exp(-\beta U(x_i)) \equiv z_i^1$

Consider a lattice (e.g. 2D-square).

At a given point x_i, the potential energy is $U(x_i)$.

The Boltzmann factor for a particle at point x_i is

$\exp(-\beta U(x_i)) \equiv z_i^1$

The partition function for a single point particle is

$Z_1 \equiv \sum_i z_i^1$

Dimers

The Boltzmann factor for a dimer on points x_i and x_{i+1} is

$\exp(-\beta(U(x_i) + U(x_{i+1})) = z_i^1 \times z_{i+1}^1$

The Boltzmann factor for all dimers terminating

on point x_i is

$z_i^{(2)} \equiv z_i^1 \times \sum_{jnni} z_j^1$

The partition function for a single dimer is

$Z_2 \equiv \sum_i z_i^{(2)}$

n-mers

The Boltzmann weight for an n-mer terminating on point x_i is

$$z_i^{(n)} = z_i^1 \times \sum_{jnni} z_j^{(n-1)}$$

and the corresponding partition function is

$$Z^{(n)} = \sum_i z_i^{(n)}$$

This method is exact for non-self-avoiding, non-interacting lattice polymers.

It can be used to speed up MC sampling of (self)interacting polymers

B. Bozorgui and DF, Phys. Rev. E 75, 036708 (2007))

NOTE: `MFOLD’ also uses recursive sampling to predict RNA secondary structures.

FREE-ENERGY METHODS OUTSIDE STATISTICAL MECHANICS

EXAMPLES:

1. Recursive analysis of Molecular Motor trajectories
2. Computation of granular entropy

FREE-ENERGY METHODS OUTSIDE STATISTICAL MECHANICS

EXAMPLES:

1. Recursive analysis of Molecular Motor trajectories
2. Computation of granular entropy

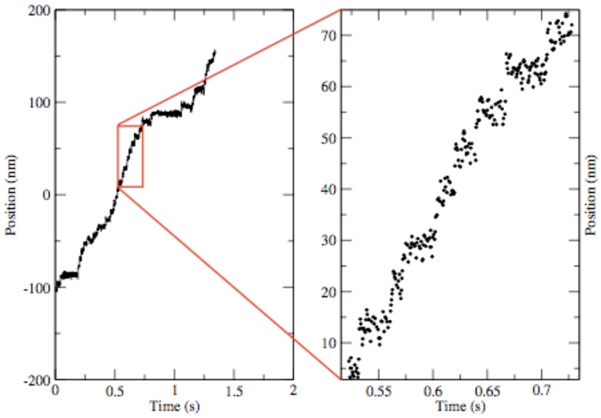

Kinesin motor steps along micro-tubules with a step size of 8nm

---

Experimentally, the step size is measured by fitting the (noisy) data.

Example: noisy “synthetic data”

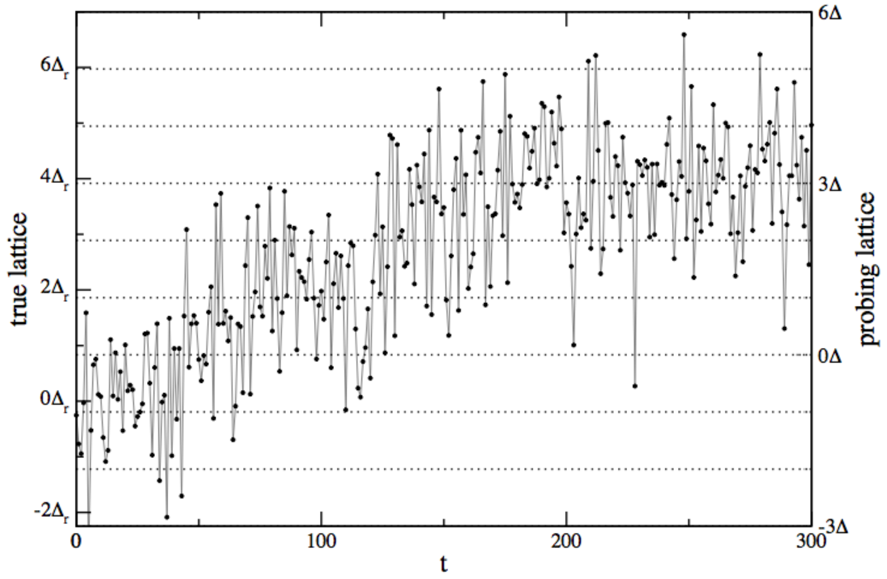

Example: noisy “synthetic data”

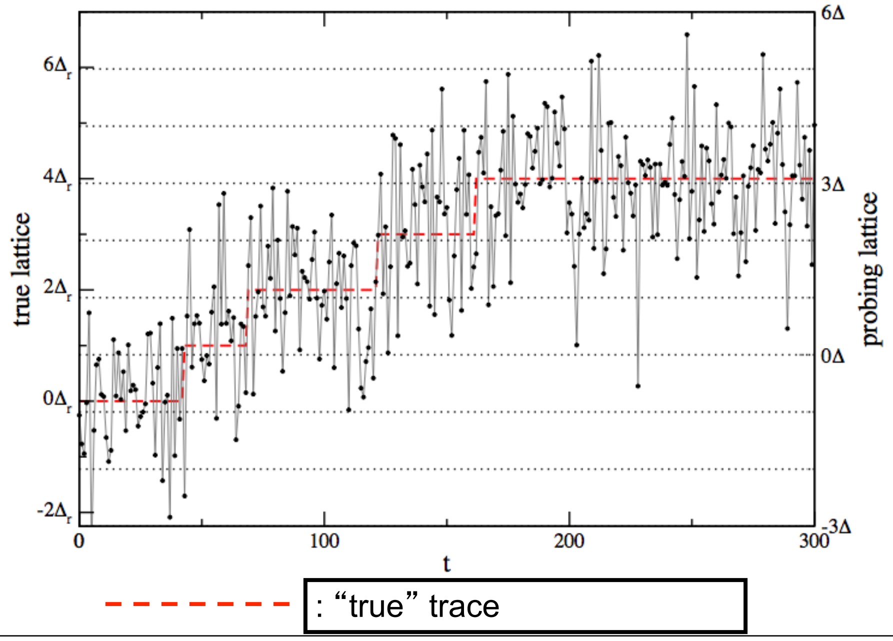

Best practice: “fit steps to data”

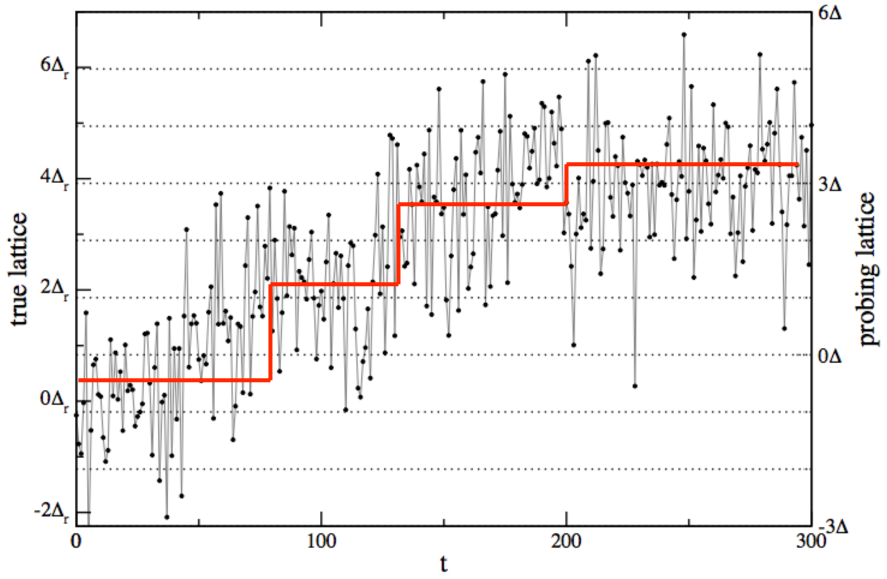

J.W.J. Kerssemakers et al. , Nature 442,709 (2006)

How well does it perform?

1. It can be used if the noise is less than 60% of the step size.
2. It yields a distribution of step sizes (even if the underlying process has only one step size)

Observation:

We want to know the step size and the step frequency but…

We do not care which trace is the “correct” trace.

Bayesian approach: compute the partition function **Q** of non-reversing polymer in a rough potential energy landscape

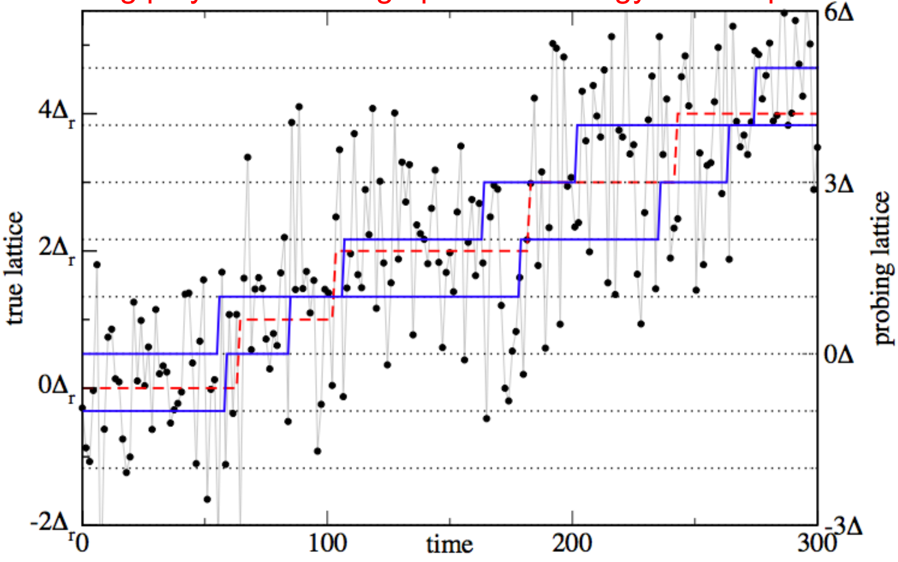

“true” trace

Other directed walks

As shown before: we can enumerate **Q** exactly (and cheaply).

From **Q** we can compute a “free energy”

$$(F = -\ln Q)$$

Compute the “excess free energy” with respect to reference data

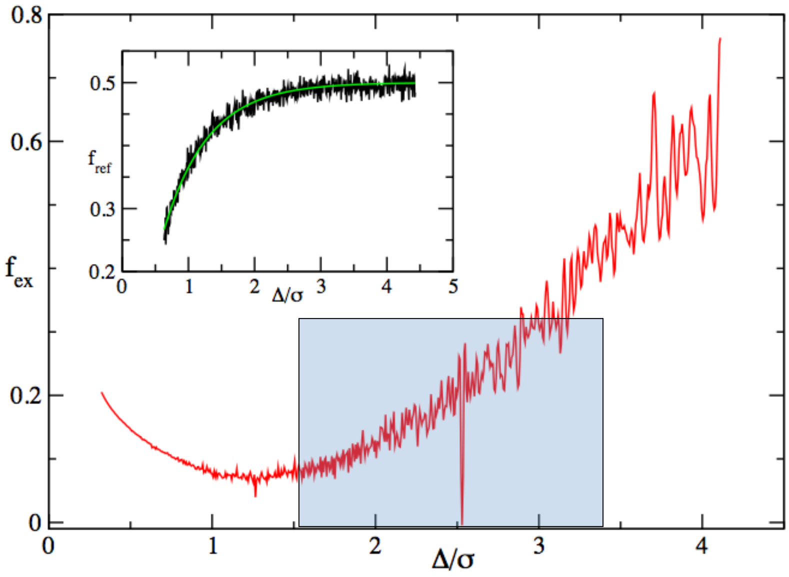

FREE-ENERGY METHODS OUTSIDE STATISTICAL MECHANICS

EXAMPLES:

1. Recursive analysis of Molecular Motor trajectories
2. Computation of granular entropy

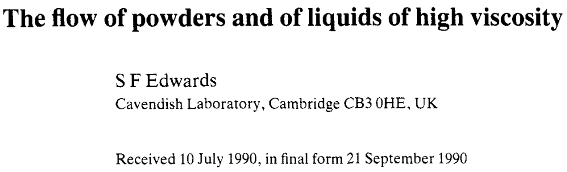

The flow of powders and of liquids of high viscosity

S F Edwards
Cavendish Laboratory, Cambridge CB3 0HE, UK

Received 10 July 1990, in final form 21 September 1990

J. Phys. Condens. Matter 2, SA63 (1990)

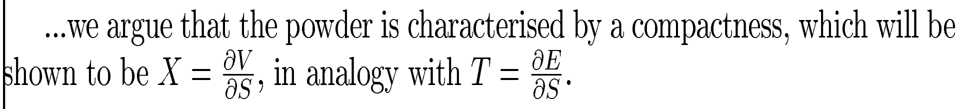

…we argue that the powder is characterised by a compactness, which will be shown to be $X = \frac{\partial V}{\partial S}$, in analogy with $T = \frac{\partial E}{\partial S}$.

`Notice that the entropy *S(N,V)* is a well defined quantity, the logarithm of the number of ways the grains can be assembled to fill the volume V…'

Well defined – maybe…

But we cannot test much, if we cannot compute $S_{granular}$

Note: `powders’ are non-thermal. Hence, Boltzmann Stat Mech does not apply.

How to count number of mechanically stable `jammed’ states?

(number of potential energy minima)

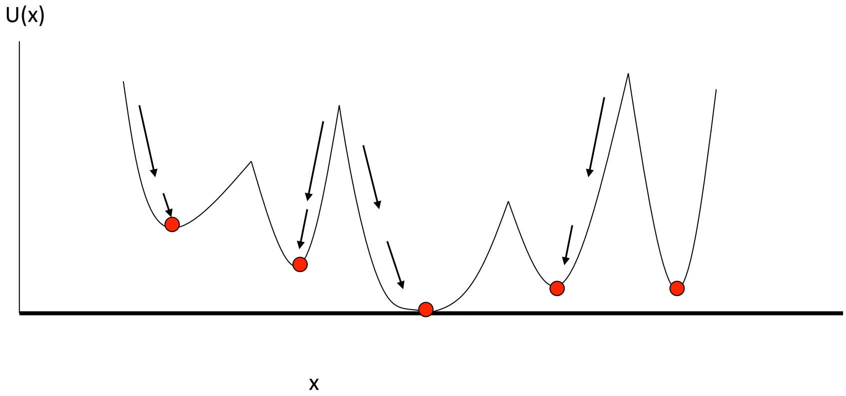

Start with a **random** initial configuration of soft spheres and

find the nearest potential-energy minimum

`High-dimensional' case

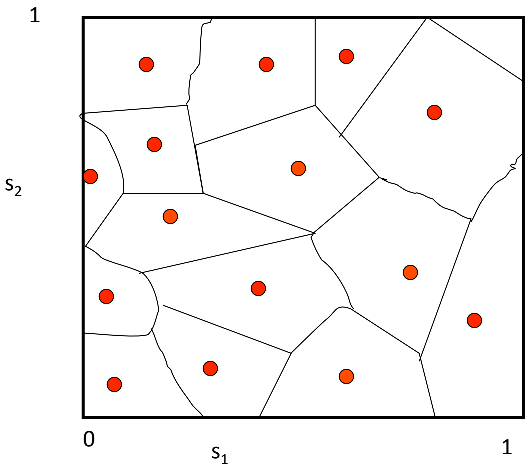

Can we count the number of distinct jammed states numerically ?

1. Brute-force method.

   Try a large number of initial configurations. Count how often a given minimum is visited.

   Works only for small systems (O(15))

2. “Average-volume” route.

Brute force method:

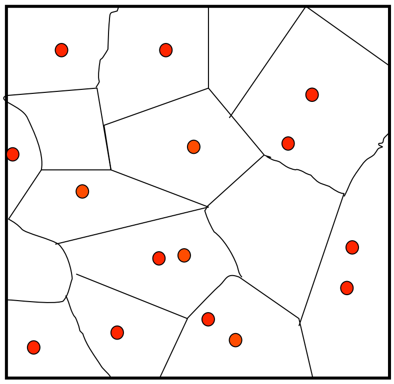

How do we count the number of distinct, disordered states?

1. Compute the distribution **P(v)** of (scaled) volumes **v.**
2. $1/N_{ds} = \langle v \rangle$

This translates a **counting** problem into a **sampling** problem.

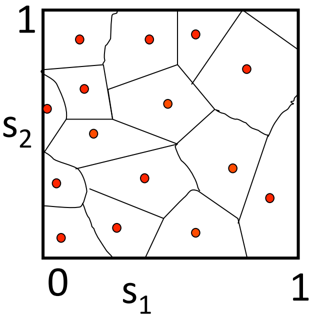

How to enumerate the troughs ("minima") in a percolating lake?

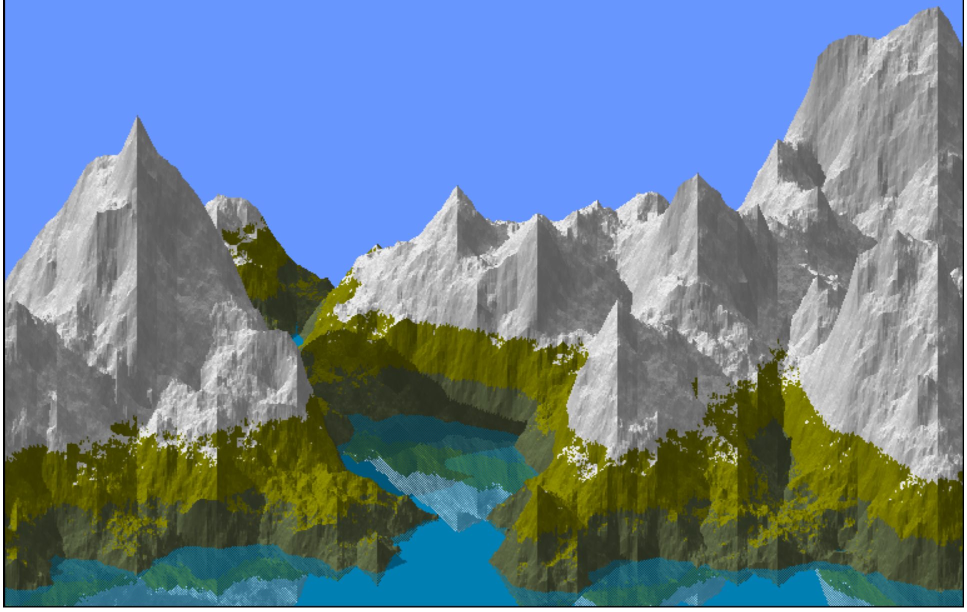

STEPS:

1. Compute the area **A** of the map (easy: L_x x L_y)
2. Compute the average area <a> of a trough ("the volume of a basin of attraction")
3. Ω = **A**/<a>

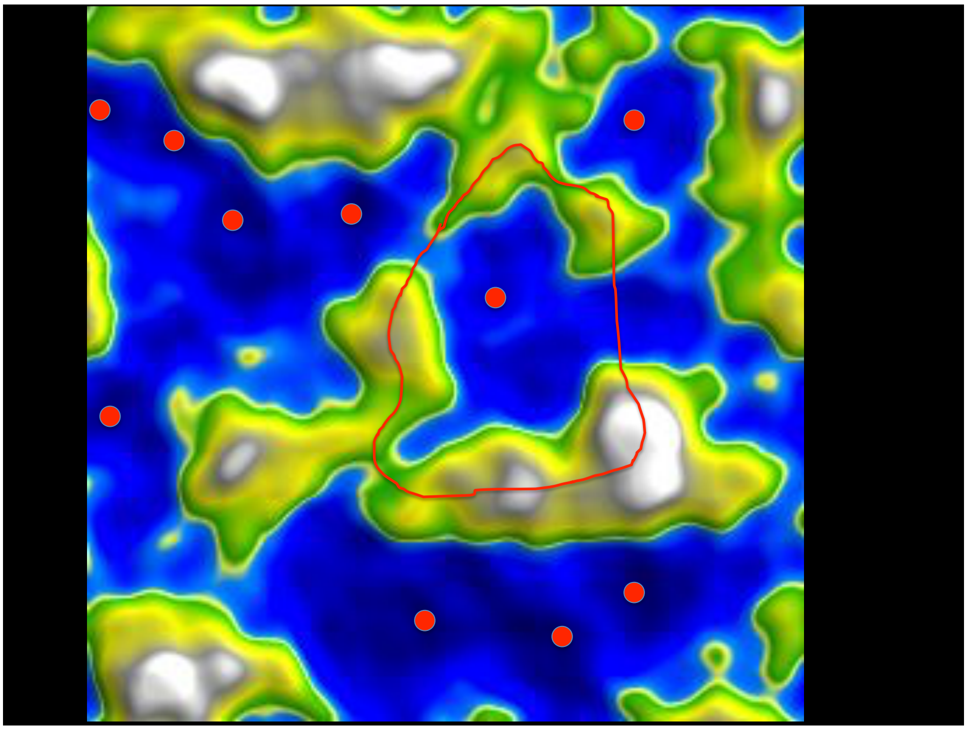

To compute the “hyper-volume” **v** of the basin of attraction of a given jammed state we must use a free-energy calculation (similar to Einstein-crystal method):

f(v)=-kT ln(**v**)

Calculation (e.g. by thermodynamic integration) is expensive because every Monte Carlo trial move requires a full energy minimization

$H_0(X) = 0$ if inside

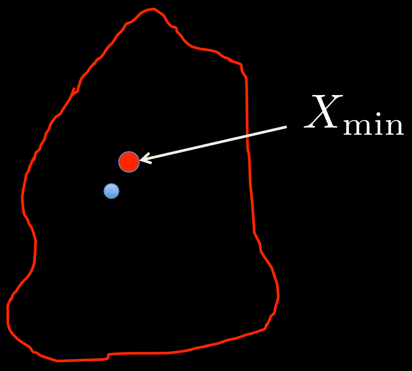

$H_0(X) = \infty$ if outside

$$V_{\text{basin}} = \int_{\text{basin}} dX \;\; \exp\left(-H_0\right)$$

Looks like a partition function…

Define `free energy' *f*

$$f = -\ln V_{\text{basin}}$$

Compute *f* by thermodynamic integration

Generalise Hamiltonian:

$$H_\lambda = H_0(X) + \lambda \left(X - X_{\min}\right)^2$$

Define `free energy' *f*(λ)

$$f(\lambda) = -\ln\left[\int_{\text{basin}} dX \;\; \exp\left(-H_\lambda(X)\right)\right]$$

For `large' λ, f(λ) is the (known) N-dimensional Harmonic Oscillator free energy.

Compute $V_{\text{basin}} = e^{-f(\lambda=0)}$

By thermodynamic integration, using

$$\frac{\partial f(\lambda)}{\partial \lambda} = \left\langle (X - X_{\text{min}})^2 \right\rangle$$

Practical challenge: we must sample **inside** the basin

Computing basin volumes in high-dimensional spaces is a general problem, not just in granular physics

Example from Dynamical Systems Theory:

"the entire topic of basins is something of an enigma in dynamical systems theory [. . .] what we do not know is how to compute the total volume or "measure" of a basin, which is what determines the probability that a random initial state will be drawn toward the associated attractor."

D. A. Wiley, S. H. Strogatz, and M. Girvan. *Chaos* 16.1 (2006), p. 015103

This is an example of the distribution of basin volumes

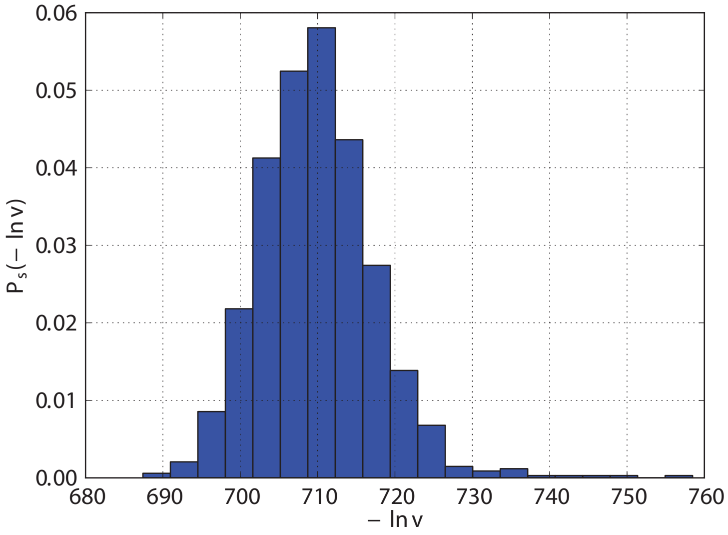

System: 2D polydispers Hard Disks

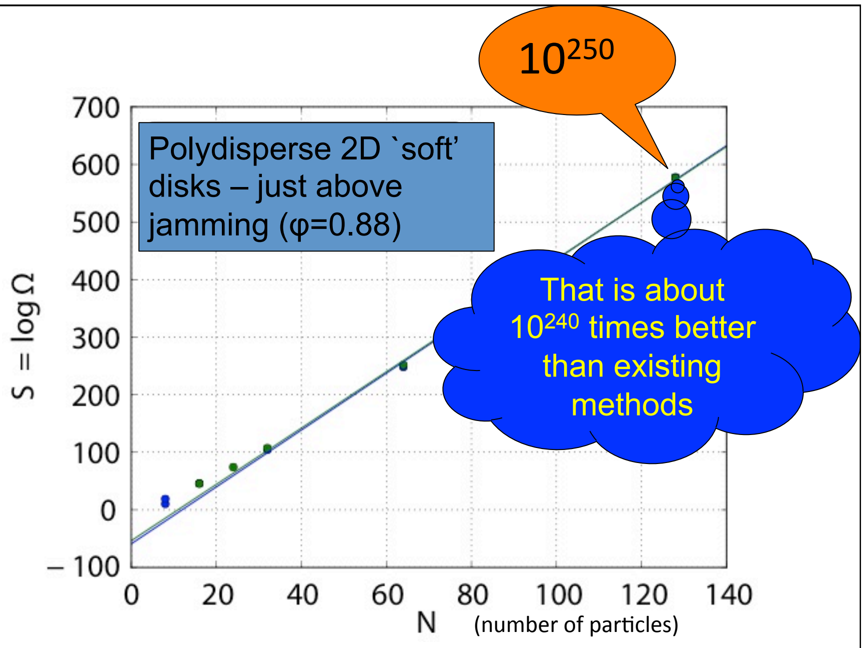

Polydisperse 2D `soft' disks – just above jamming (φ=0.88)

10^{250}

That is about 10^{240} times better than existing methods

MBAR

(Multistate Bennett Acceptance Ratio)

Method to obtain the best estimate of free-energy differences from umbrella sampling

Shirts, M. R., and Chodera, J. D. (2008) Statistically optimal analysis of samples from multiple equilibrium states. *J. Chem. Phys.* 129, 129105.

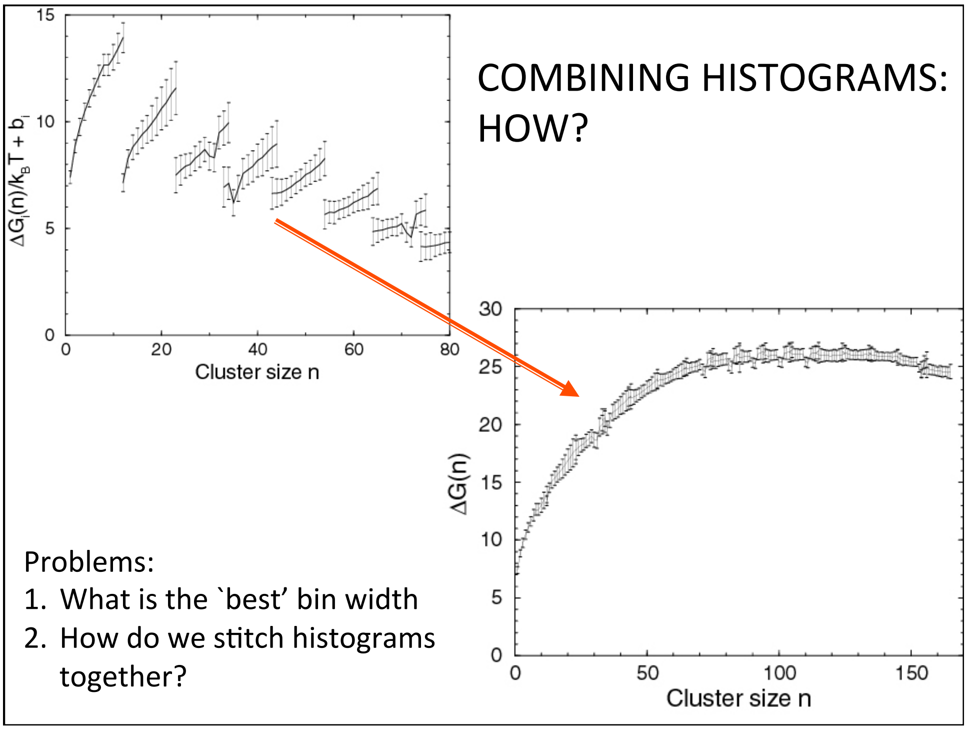

COMBINING HISTOGRAMS: HOW?

Problems:

1. What is the `best' bin width
2. How do we stitch histograms together?

MBAR: No binning and `optimal' stitching.

We start from:

$$Z = \int d\mathbf{R}^N \ \exp[-\beta U(\mathbf{R}^N)]$$

and

$$F = -k_B T \ln Z$$

Suppose we have k different samples (e.g. in umbrella sampling), biased with potentials $V_k(\mathbf{R}^N)$. Assume that we have N_k points for sample k We can then define ‘partition functions Z_k for the biased systems as

$$Z_k \equiv \int d\mathbf{R}^N \exp(-\beta[U(\mathbf{R}^N) + V_k(\mathbf{R}^N)])$$

and

$$F_k \equiv -k_B T \ln Z_k$$

In what follows, we will use:

$$\Delta F_k \equiv F_k - F = k_B T \ln(Z/Z_k)$$

The key assumption of MBAR is that the true (as opposed to the sampled) distribution function is a weighted set of delta-functions *at the points that have been sampled*.

In words: we do not assume anything about points that we have not sampled.

The distribution function is then of the form:

$$P(\mathbf{R}^N) = \mathcal{Z}^{-1} \sum_{j=1}^{K} \sum_{n=1}^{N_k} p_{j,n} \delta\left(\mathbf{R}^N - \mathbf{R}^N_{j,n}\right)$$

Where the $p_{j,n}$ are (as yet) unknown.

The normalization factor is defined as:

$$\mathcal{Z} \equiv \sum_{j=1}^{K} \sum_{n=1}^{N_k} p_{j,n}$$

Once the full distribution is known, the biased distributions follow:

$$P_k(\mathbf{R}^N) = \mathcal{Z}_k^{-1} \sum_{j=1}^{K} \sum_{n=1}^{N_k} p_{j,n} \exp(-\beta V_k(\mathbf{R}^N)) \delta\left(\mathbf{R}^N - \mathbf{R}_{j,n}^N\right)$$

The normalization factor Z_k is defined as:

$$\mathcal{Z}_k \equiv \sum_{j=1}^{K} \sum_{n=1}^{N_k} p_{j,n} \exp(-\beta V_k(\mathbf{R}_{j,n}^N))$$

Now we must compute the unknown weights $p_{j,n}$

We do this, using `maximum likelihood’.

That is: we impose that the values of the $p_{j,n}$ should be such that the probability of obtaining the observed histograms is maximised

We define the *likelihood* **L:**

$$L \equiv \prod_{j=1}^{K} \left[\prod_{n=1}^{N_k} P_k(\mathbf{R}_{j,n}^N)) \right]$$

L depends on all $p_{j,n}$

We determine $p_{j,n}$ by imposing that **L**, or equivalently ln **L** is maximal.

If we look at ln **L**

$$\ln L \equiv \sum_{j=1}^{K} \sum_{n=1}^{N_k} \ln \left[\frac{p_{j,n}}{\mathcal{Z}_k} \exp(-\beta V_k(\mathbf{R}_{j,n}^N)) \right]$$

We see that ln $p_{j,n}$ and Z_k depend on $p_{j,n}$
But the Boltzmann factor does not.

Therefore:

$$\begin{aligned}\ln L &= \text{constant} + \sum_{j=1}^{K}\sum_{n=1}^{N_k}\left[\ln p_{j,n} - \ln \mathcal{Z}_j\right]\\ &= \text{constant} + \sum_{j=1}^{K}\sum_{n=1}^{N_k}\ln p_{j,n} - \sum_{j=1}^{K} N_j \ln \mathcal{Z}_j\end{aligned}$$

Now, we can differentiate with respect to $p_{j,n}$

The constant yields zero.

The second term: $1/p_{j,n}$

The third term follows if we use:

$$\mathcal{Z}_k \equiv \sum_{j=1}^{K}\sum_{n=1}^{N_k} p_{j,n} \exp(-\beta V_k(\mathbf{R}_{j,n}^N))$$

Our condition for maximum likelihood is then

$$0 = \frac{1}{p_{j,n}} - \sum_{k=1}^{K} N_k \frac{\exp[-\beta V_k(\mathbf{R}_{j,n}^N))]}{\mathcal{Z}_k}$$

Or:

$$p_{j,n}/\mathcal{Z} = \frac{1}{\sum_{k=1}^{K} N_k \frac{\exp[-\beta V_k(\mathbf{R}_{j,n}^N))]}{(\mathcal{Z}_k/\mathcal{Z})}}$$

The probability to observe a given point (j,n) given the optimal $p_{j,n}$ is then

$$p_{j,n}/\mathcal{Z} = \frac{1}{\sum_{k=1}^{K} N_k \exp[-\beta(V_k(\mathbf{R}_{j,n}^N) - \Delta F_k)]}$$

Where we have used

$$\Delta F_k \equiv k_B T \ln(Z/Z_k) \approx k_B T \ln(\mathcal{Z}/\mathcal{Z}_k)$$

We can rewrite our result as an implicit equation for the ΔF_i :

$$\Delta F_i = -k_B T \ln \sum_{j=1}^{K} \sum_{n=1}^{N_j} \frac{\exp[-\beta(V_i(\mathbf{R}_{j,n}^N)]}{\sum_{k=1}^{K} N_k \exp[-\beta(V_k(\mathbf{R}_{j,n}^N) - \Delta F_k)]}$$

These are the MBAR equations that must be solved self-consistently

Advantages of MBAR over all earlier schemes (except Bennett)

- It does not use bins.
- it makes no assumption about the form of the distribution function where it has not been sampled.
- different biased runs may sample different points in parameter space
- the method yields the *best* (in the sense of `the most likely') estimate for the histograms and the free energy differences.