Introduction to Macromolecular Simulation

Peter J. Steinbach

Center for Molecular Modeling
Center for Information Technology
National Institutes of Health
Bldg. 12A Room 2051
steinbac@helix.nih.gov

To print this tutorial on a PostScript printer, download intro_simulation.ps.

Classical Mechanics Applied to Biology

The purpose of this tutorial is to introduce several popular numerical techniques used to simulate the structure and dynamics of biomolecules. The discussion is confined to simulation methods that apply classical mechanics to biological systems, although some quantum theory is presented to quantify some shortcomings of classical approximations. Molecular dynamics (MD) simulation, Langevin dynamics (LD) simulation, Monte Carlo (MC) simulation, and normal mode analysis are among the methods surveyed here. There are techniques being developed that treat the bulk of a macromolecule classically while applying quantum mechanics to a subset of atoms, typically the active site. This research frontier will not be addressed here. Completely classical studies remain more common and continue to contribute to our understanding of biological systems.

When is classical mechanics a reasonable approximation?

In Newtonian physics, any particle may possess any one of a continuum of energy values. In quantum physics, the energy is *quantized*, not continuous. That is, the system can accomodate only certain discrete levels of energy, separated by gaps. At very low temperatures these gaps are much larger than thermal energy, and the system is confined to one or just a few of the low-energy states. Here, we expect the `discreteness' of the quantum energy landscape to be evident in the system's behavior. As the temperature is increased, more and more states become thermally accessible, the `discreteness' becomes less and less important, and the system approaches classical behavior.

For a harmonic oscillator, the quantized energies are separated by $\Delta E = hf$, where h is Planck's constant and f is the frequency of harmonic vibration. Classical behavior is approached at temperatures for which $k_B T \gg hf$, where k_B is the Boltzmann constant and $k_B T$ = 0.6 kcal/mol at 300 K.

Setting hf = 0.6 kcal/mol yields f = 6.25/ps, or 206 cm^{-1}. So a classical treatment will suffice for motions with characteristic times of a ps or longer at room temperature.

Outline - Shades of things to come

We'll expand on the above argument with a more quantitative analysis of classical and quantum treatments of simple harmonic oscillation. This not-too-mathematical glimpse of quantum mechanical

phenomena is included to help simulators estimate how much they can trust various motions that have been simulated with the approximations inherent in classical physics. Then, we'll identify the basic ingredients of a macromolecular simulation: a description of the structure, a set of atomic coordinates, and an empirical energy function. This is followed by a discussion of the most popular simulation techniques: energy minimization, molecular dynamics and Monte Carlo simulation, simulated annealing, and normal-mode analysis. Finally, a few general suggestions are offered to those about to perform their first macromolecular simulation. But first, a little theoretical background is presented to aid the discussion. It's a short summary of the most relevant concepts of classical, quantum, and statistical mechanics, along with a glimpse of classical electrostatics.

- Classical and Quantum Mechanics - in a Nutshell
- Statistical Mechanics - Calculating Equilibrium Averages
- Classical vs. Quantum Mechanics: The Harmonic Oscillator in One Dimension
    - Probability(x): Where is the oscillating particle?
    - Average Energy U and Heat Capacity
    - Mean-Square Fluctuation
    - Overall Comparison - What does all this mean for simulation?

- Electrostatics and the `Generalized Born' Solvent Model
- Classical Macromolecular Simulation
    - A Note on Notation - Some CHARMMing language
    - Simulating Biomolecules: The Three Necessary Ingredients
        - I. A Description of the Structure
        - II. Initial Coordinates
   - III. An Empirical Energy Function: Free Energy vs. Potential Energy

- The Empirical Potential Energy Function
    - Matching CHARMM's Electrostatic Approximations to Environmental Approximations
        - No Explicit Solvent (RDIE)
        - Partial Solvation by Explicit Water Molecules
        - Solutions, Crystals, and Interfaces

- Classical Simulation and Modeling Techniques
    - Questions We Can Ask With a Computer
    - Energy Minimization
    - Molecular Dynamics (MD) Simulation
    - Langevin Dynamics (LD) Simulation
    - Monte Carlo (MC) Simulation
    - Normal Mode (Harmonic) Analysis
    - Simulated Annealing

- What is Unique to Computer Experiments?
- Worth Worrying About:
- Bibliography
- About this document ...

Next: [Classical and Quantum Mechanics](#)
Steinbach 2003-03-07

Classical and Quantum Mechanics - in a Nutshell

Classical Mechanics

Building on the work of Galileo and others, Newton unveiled his laws of motion in 1686. According to Newton:

I. A body remains at rest or in uniform motion (constant velocity - both speed and direction) unless acted on by a net external force.

II. In response to a net external force, $\vec{F}$, a body of mass m accelerates with acceleration $\vec{a} = \vec{F}/m$.

III. If body i pushes on body j with a force $\vec{F}_{ij}$, then body j pushes on body i with a force

$$\vec{F}_{ji} = -\vec{F}_{ij}.$$

For energy-conserving forces, the net force $\vec{F}_i$ on particle i is the negative gradient (slope in three dimensions) of the potential energy with respect to particle i's position: $\vec{F}_i = -\vec{\nabla}_i V(\vec{R})$, where

$V(\vec{R})$ represents the potential energy of the system as a function of the positions of all N particles, $\vec{R}$.

In three dimensions, $\vec{r}_i$ is the vector of length 3 specifying the position of the i^{th} atom, and $\vec{R}$ is the vector of length $3N$ specifying all coordinates. In the context of simulation, the forces are calculated for energy minimizations and molecular dynamics simulations but are not needed in Monte Carlo simulations.

Classical mechanics is completely deterministic: Given the exact positions and velocities of all particles at a given time, along with the function $V(\vec{R})$, one can calculate the future (and past) positions and velocities of all particles at any other time. The evolution of the system's positions and momenta through time is often referred to as a *trajectory*.

Quantum Mechanics

A number of experimental observations in the late 1800's and early 1900's forced physicists to look beyond Newton's laws of motion for a more general theory. See, for example, the discussion of the heat capacity of solids. It had become increasingly clear that electromagnetic radiation had particle-like

properties in addition to its wave-like properties such as diffraction and interference. Planck showed in 1900 that electromagnetic radiation was emitted and absorbed from a black body in discrete quanta, each having energy proportional to the frequency of radiaion. In 1904, Einstein invoked these quanta to explain the photo-electric effect. So under certain circumstances, one must interpret electromagnetic waves as being made up of particles. In 1924 de Broglie asserted that matter also had this dual nature: Particles can be wavey.

To make a long and amazing story [1] short, this led to the formulation of Shrödinger's wave equation for matter:

$$H\Psi = E\Psi.$$

Don't let the brevity of notation fool you; this partial differential equation is difficult to deal with and generally impossible to solve analytically. It is tailored to a given physical system by defining the *Hamiltonian operator* H to incorporate all the relevant forces exerted on the particles of the system. The solution of this equation yields the discrete (*quantized*) values (or *eigenvalues*) of energy E_n, and for each E_n its corresponding *wave function* Ψ_n. In general, these wave functions are *complex*-valued functions (involving $\sqrt{-1}$), but the quantity $\Psi^*\Psi$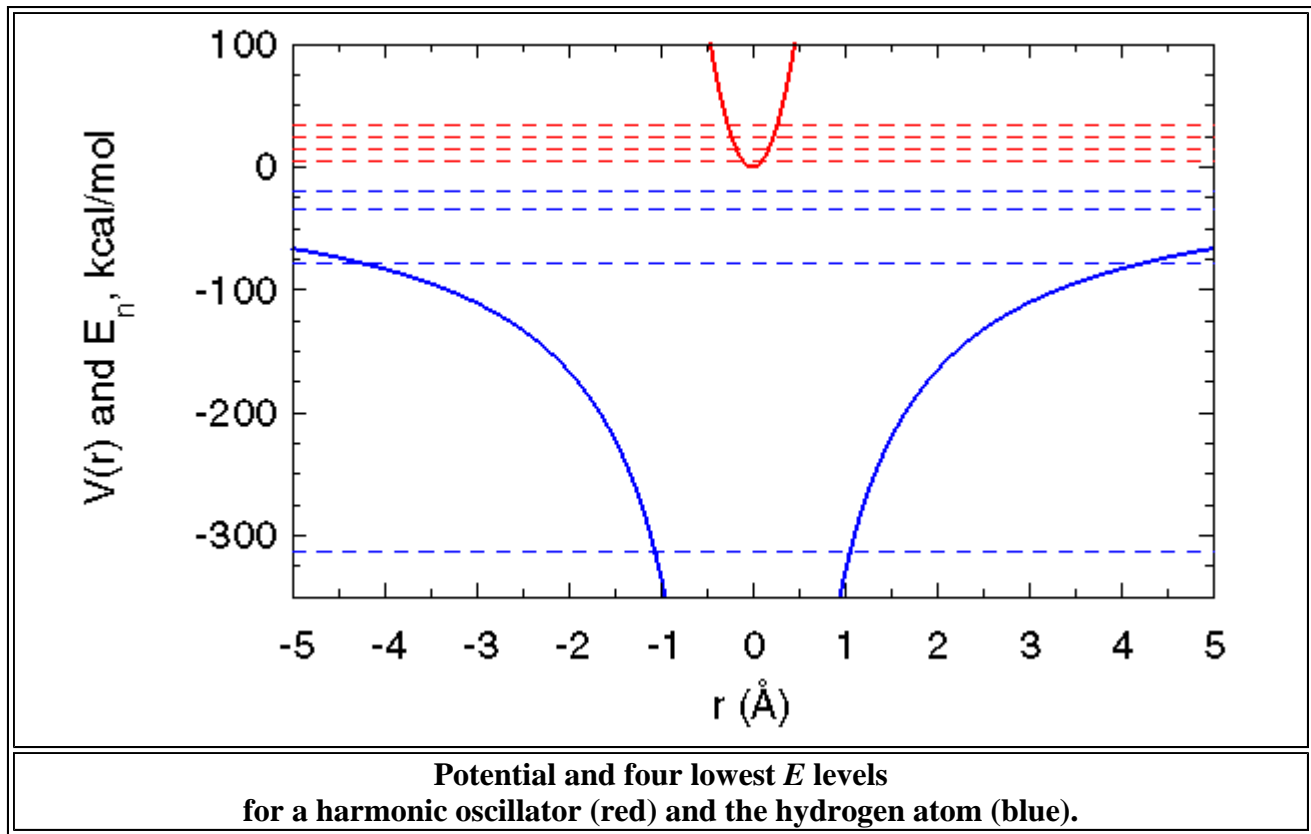 is always *real* and thus may correspond to something physical. (Ψ^* is the `complex conjugate' of $\Psi$.) In fact, $\Psi^*\Psi$ is a *probability density*. For motion in the single dimension *x*, it is `a probability per unit *x*': $\Psi^*\Psi dx$ is the probability that the particle will be found at a position between *x* and $x + dx$. The wavefunctions are *normalized* (scaled) by the requirement that the particle must be somewhere, i.e., that these probabilities must sum to one:

$$\int \Psi^*\Psi \vec{dr} = 1.$$

Quantum mechanics is thus not deterministic, but probabilistic. It forces us to abandon the notion of precisely defined trajectories of particles through time and space. Instead, we must talk in terms of *probabilities* for alternative system configurations.

To clarify these concepts, consider two major successes for the quantum theory, predictions of the discrete energy levels of the harmonic oscillator and the hydrogen atom. Pictured below are the potential energy (solid lines) and the four lowest energy levels (dashed lines) for a one dimensional harmonic oscillator (red) and the three dimensional hydrogen atom (blue). The harmonic oscillator depicted corresponds to a hydrogen atom oscillating at the frequency *f* = 100/ps and represents one of the highest frequency atomic motions in macromolecules. The energy levels of harmonic oscillators are equally spaced, separated by an energy of *hf*, or 9.5 kcal/mol for the oscillator shown. The energy gaps for a hydrogen atom oscillating at *f* = 10/ps are 0.95 kcal/mol, on the order of thermal energy, and so classical mechanics better approximates quantum results (e.g., average energy and motional amplitude) for this slower oscillator.

Excitation of electrons within atoms requires much more energy than excitation of atomic vibrations. Promotion of the hydrogen atom's electron from its ground state to its first excited state requires 235 kcal/mol. Way beyond the reach of thermal energy, this excitation requires the absorption of ultraviolet radiation with a wavelength of 121 nm.

Potential and four lowest *E* levels
for a harmonic oscillator (red) and the hydrogen atom (blue).

Steinbach 2003-03-07

Statistical Mechanics - Calculating Equilibrium Averages

According to statistical mechanics, the probability that a given state with energy E is occupied in equilibrium at constant particle number N, volume $\mathcal{V}$, and temperature T (constant $N\mathcal{V}T$, the `canonical' ensemble) is proportional to $e^{-E/k_BT}$, the `Boltzmann factor.'

$$probability \propto e^{-E/k_BT}. \tag{1}$$

The equilibrium value of any observable O is therefore obtained by averaging over all states accessible to the system, weighting each state by this factor.

Quantum mechanically, this averaging is performed simply by *summing* over the *discrete* set of states (Figure 1):

$$\langle O \rangle = \frac{\sum_n O_n e^{-E_n/k_BT}}{Z}, \tag{2}$$

where Z is the *partition function*:

$$Z = \sum_n e^{-E_n/k_BT}, \tag{3}$$

and O_n is the *expectation value* of the quantity O in the n^{th} energy eigenstate:

$$O_n = \int \Psi_n^* O \Psi_n \vec{dr}. \tag{4}$$

Classically, a state is specified by the positions and velocities (momenta) of all particles, each of which can take on any value. The averaging over states in the classical limit is done by *integrating* over these *continuous* variables:

$$\langle O \rangle = \frac{\int O e^{-E/k_BT} \vec{dp}\vec{dr}}{\int e^{-E/k_BT} \vec{dp}\vec{dr}},$$

where the integrals are over all phase space (positions $\vec{r}$ and momenta $\vec{p}$) for the N particles in 3 dimensions.

When all forces (the potential energy *V*) and the observable *O* are velocity-independent, the momentum integrals can be factored and canceled:

$$\langle O \rangle = \frac{\int e^{-K/k_B T} \vec{dp} \int O e^{-V/k_B T} \vec{dr}}{\int e^{-K/k_B T} \vec{dp} \int e^{-V/k_B T} \vec{dr}} = \frac{\int O e^{-V/k_B T} \vec{dr}}{\int e^{-V/k_B T} \vec{dr}}, \tag{5}$$

where $K = \sum_{i=1}^{N} p_i^2 / 2m_i$ is the total kinetic energy, and $E = K + V$. As a result, Monte Carlo

simulations compare *V*'s, not *E*'s.

Steinbach 2003-03-07

Classical vs. Quantum Mechanics: The Harmonic Oscillator in One Dimension

The harmonic oscillator is the model system of model systems. We study it here to characterize differences in the dynamical behavior predicted by classical and quantum mechanics, stressing concepts and results. More details and mathematical formalism can be found in textbooks [1,2]. Our model system is a single particle moving in the *x* dimension connected by a spring to a fixed point. Its potential energy is $V(x) = kx^2/2$, where *k* is the spring constant. Stiff springs are described by large *k*'s.

Classically, this oscillator undergoes sinusoidal oscillation of amplitude $A = \sqrt{2E/k}$ and frequency

$f = (k/m)^{1/2}/(2\pi)$, where *E* is the total energy, potential plus kinetic. In equilibrium at temperture

T, its average potential energy and kinetic energy are both equal to $k_BT/2$; they depend only on

temperature, not on the motion's frequency.

Quantum mechanically, the *probability* of finding the particle at a given place is obtained from the solution of Shrödinger's equation, yielding eigenvalues E_n and eigenfunctions $\Psi_n(x)$. For the one

dimensional harmonic oscillator, the energies are found to be $E_n = (n + 1/2)hf$, where h is

Planck's constant, *f* is the classical frequency of motion (above), and *n* may take on integer values from 0 to infinity. The $\Psi_n(x)$ turn out to be real functions involving the *Hermite polynomials*. From

equation 1, only the *ground state* ($n = 0$) is populated as the temperature $T \rightarrow 0$. The energy does not go to zero but to $hf/2$. The corresponding *zero-point motion* is a quantum mechanical phenomenon.

Classically, there is no motion as $T \rightarrow 0$. Thus, we expect that quantum mechanics predicts more motion than classical mechanics, especially at low temperature.

Subsections

- Probability(x): Where is the oscillating particle?
- Average Energy *U* and Heat Capacity
- Mean-Square Fluctuation
- Overall Comparison - What does all this mean for simulation?

---

   previc

**Next:** Probability(x): Where is the **Up:** intro_simulation **Previous:** Statistical Mechanics - Calculating
Steinbach 2003-03-07

# Probability(x): Where is the oscillating particle?

For the oscillator in the n^{th} eigenstate with energy E_n, the probability of being between x and

$x + dx$ is $\Psi_n^* \Psi_n dx$. In the following figure, $\Psi_n^* \Psi_n$ is plotted in solid lines for $n = 0$ (red) and for

$n = 5$ (blue), along with the corresponding classical predictions, plotted in dotted lines for classical oscillators with the same total energies (0.5 hf and 5.5 hf). Numerical constants chosen: m = mass of H atom, f = 100/ps, about the frequency of O-H bond stretching.

Quantum and classical probability of finding oscillator as a function of position.
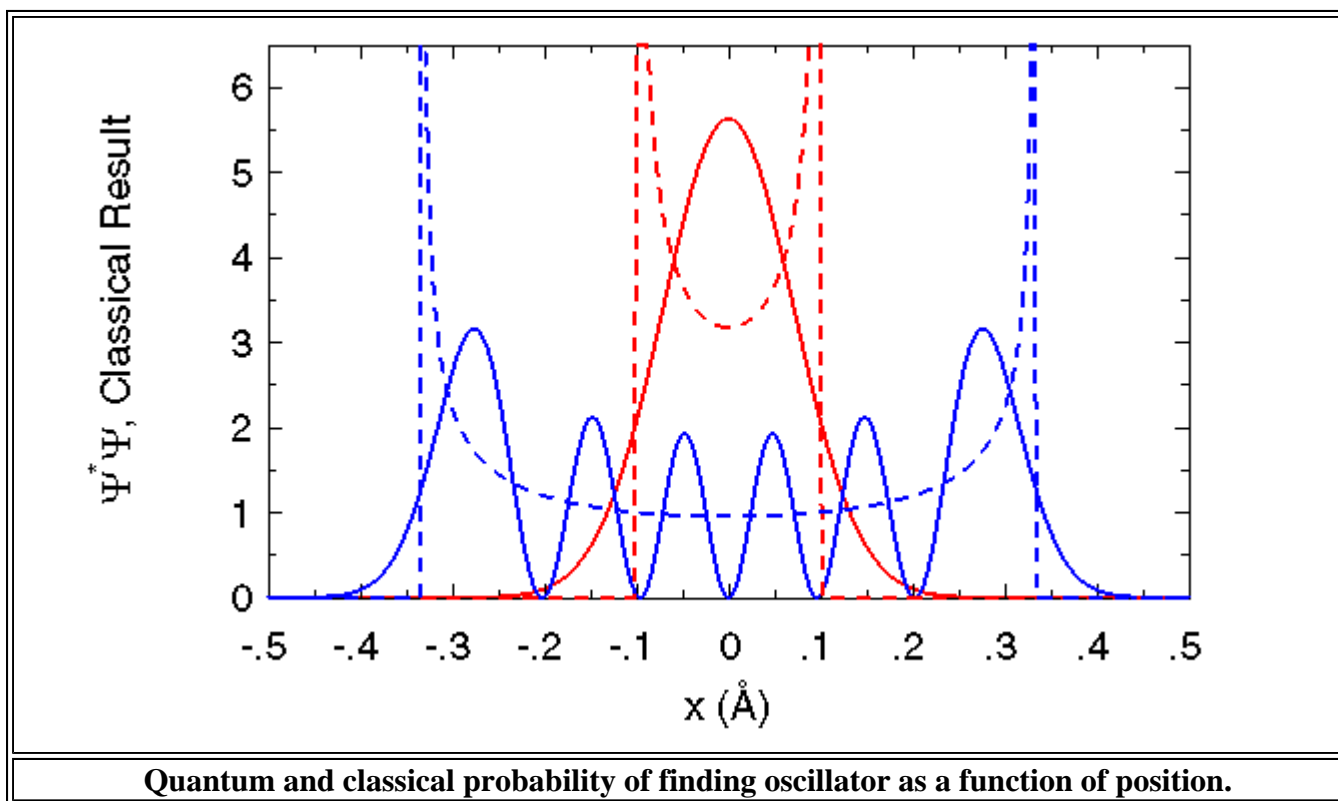
Classically, the probability that the oscillating particle is at a given value of x is simply the fraction of time that it spends there, which is inversely proportional to its velocity $v(x)$ at that position. The

particle must stop completely (for a moment) before reversing its direction, and so it spends the most time where the spring is either fully compressed or fully extended ($x = \pm A$). It spends the least time

where its velocity is greatest, i.e., where the spring is at its equilibrium length ($x = 0$). Classically, there is zero chance for a particle to have a potential energy V greater than its total energy E, and so the motion is strictly confined to the range $-A \leq x \leq A$ (see vertical dotted lines at $x = \pm A$).

Quantum mechanically, there exist states (any *n* > 0) for which there are locations *x*, where the probability of finding the particle is zero, and that these locations separate regions of high probability! Also, note that there is appreciable probability that the particle can be found outside the range $-A \leq x \leq A$, where classically it is strictly forbidden! This *quantum mechanical tunneling* is related to the famous *Heisenberg Uncertainty Principle*, which states that one *cannot know both* the position and momentum of a particle with infinite precision at the same time.

Finally, note that the classical approximation more closely resembles the average of $\Psi_n^* \Psi_n$ as the energy E_n of the oscillator increases.

Steinbach 2003-03-07

Average Energy *U* and Heat Capacity

Next, we consider the average energy $U = \langle E \rangle$ of the harmonic oscillator at fixed $N \mathcal{V} T$. We

calculate it as a function of temperature using equations 2, 4, and 5, with the observable $O = E$. *U* is plotted in dotted lines for three frequencies (spring constants *k*): $f = 100$/ps (green), 10/ps (blue), and 1/ps (cyan). The classical result (red) is $U = k_B T$, regardless of the frequency of the harmonic motion

(the classical `equipartition of energy'). Note that the deviation of classical from quantum behavior is reduced as the frequency, and hence the energy gap between quantum states, is reduced.

Harmonic oscillation: average energy and heat capacity versus temperature
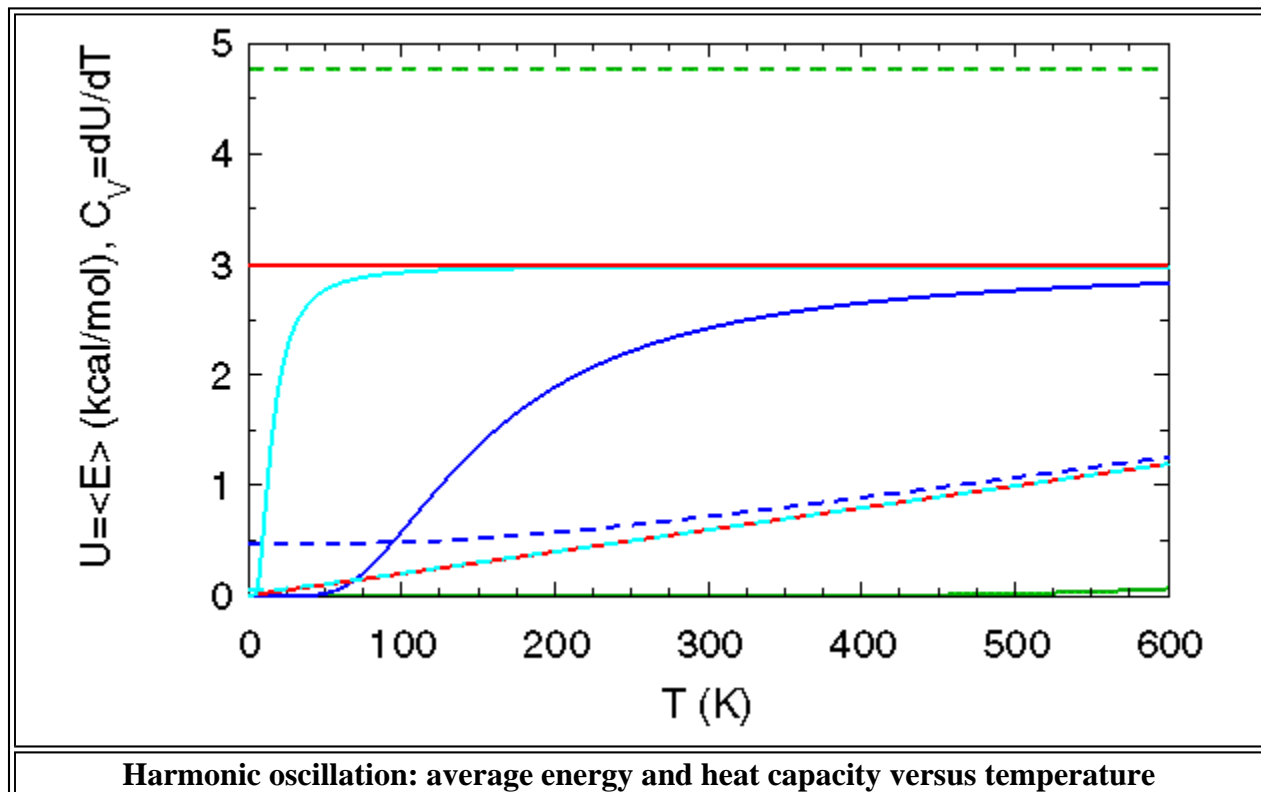
The *heat capacity* at constant volume, $C_V = \partial U / \partial T |_V$, is plotted in solid lines for each of the four *U*

curves. (To fit on the same scale, C_V values were scaled by a factor of 1500). This heat capacity

measures the energy absorbed (released) by a system as its temperature is raised (lowered) one degree while the volume is held constant. (The heat capacity *at constant pressure*, C_P, is larger than C_V for

gases because work is done as the container expands.) In the early twentieth century, experimentally measured values of the heat capacity for solids at low temperature, which can be approximated as three dimensional arrays of harmonic oscillators, pointed to a problem in the classical theory. Classical

physics predicted a temperature independent heat capacity (red horizontal line), whereas the measured values went to zero as T^3 as $T \to 0$. The much improved predictions (by Einstein in 1907 and Debye in 1912) of the low-temperature heat capacity of solids were among the earliest successes for theoretical models that invoked quantum concepts.

Harmonic oscillation at 300 K: average energy and heat capacity versus frequency

The frequency dependencies of U and C_V at 300 K are shown in the second figure. Again, the C_V values have been scaled by a factor of 1500. Classical quantities (in red) are independent of frequency, with $U = k_B T$ = 0.6 kcal/mol drawn as a dashed line. Quantum mechanical quantities (U in green, C_V in cyan) deviate more and more from classical values as the frequency is increased. As foreshadowed in the introduction, this deviation becomes appreciable at frequencies above 1/ps.

An aside: For those unfamiliar with the concept of heat capacity, consider the well known consequences of water's large heat capacity. Water can absorb and release considerable thermal energy with little change in temperature. Hence, average coastal temperatures are cooler in summer than inland temperatures as the water absorbs and stores heat. The reverse is true in winter, as the water releases the stored energy.

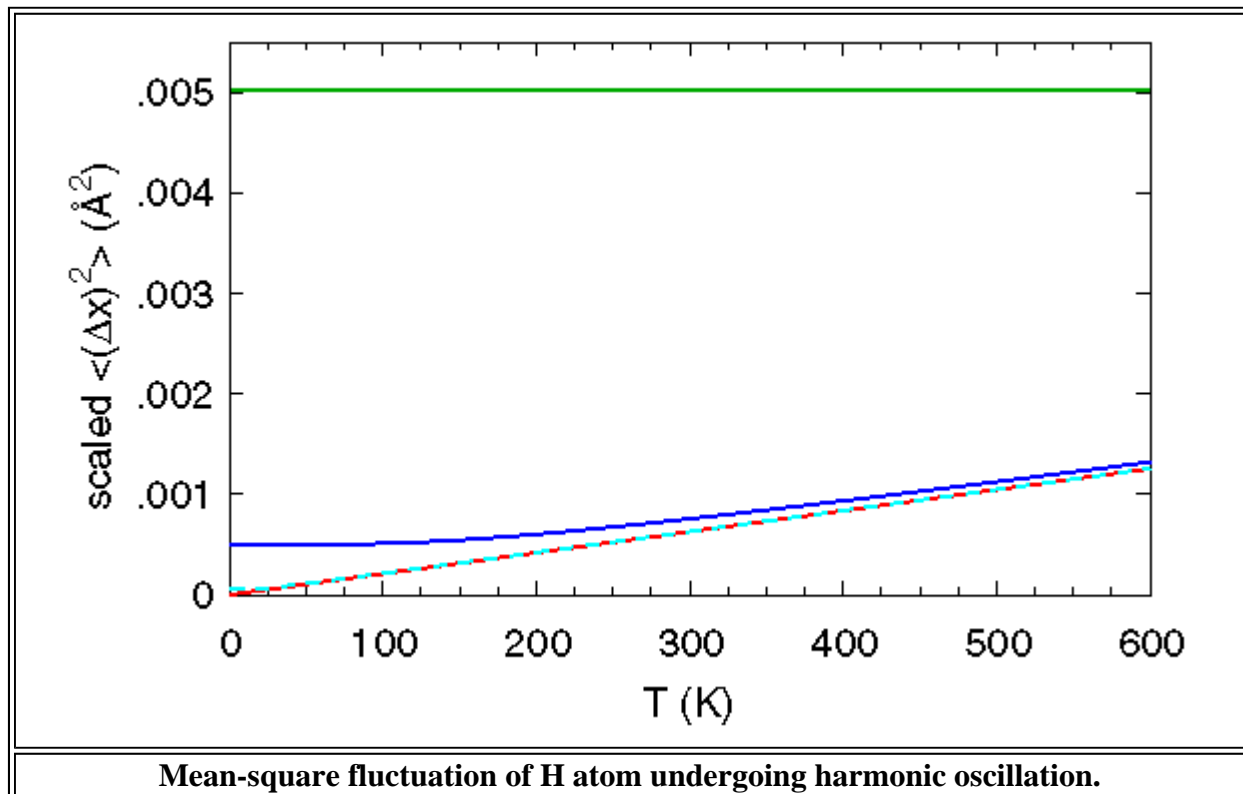*Steinbach 2003-03-07*

Mean-Square Fluctuation

One physical quantity of great interest is the variance in the position of atoms at equilibrium, $\langle(\Delta x)^2\rangle$.

For our model oscillator, $\langle x\rangle = 0$; so $\langle(\Delta x)^2\rangle = \langle x^2\rangle$. This mean-square fluctuation about the

average position is related to the *B factors* of crystallography and is also measurable by neutron scattering [3] and by Mössbauer spectroscopy [4]. It is one of the most important quantities to keep an eye on in molecular dynamics simulations as well. What is this fluctuation for the harmonic oscillator in equilibrium at constant NVT according to classical and quantum mechanics? We again use equations 2, 4, and 5, now with $O = x^2$, and consider the same three frequencies of proton vibration. Because $V = kx^2/2$ and $\langle V\rangle = \langle K\rangle = \langle E\rangle/2$ for harmonic oscillation, the quantum and classical

results are proportional to those obtained above for the average energy. That is, $\langle x^2\rangle = U/k = U/m(2\pi f)^2$. Again, *m* is taken as the mass of a proton. (To plot the three frequencies

on one scale, results have been scaled to the $f = 100/ps$ values (green): Results for $f = 10/ps$ (blue) were scaled by 0.01 ; results for $f = 1/ps$ (cyan) were scaled by 0.0001. Consequently, the three classical curves (red), $\langle x^2\rangle = k_B T/k$, coincide.)

Mean-square fluctuation of H atom undergoing harmonic oscillation.

---

   previc

Next: Overall Comparison - What **Up:** Classical vs. Quantum Mechanics: **Previous:** Average Energy U and
Steinbach 2003-03-07

Overall Comparison - What does all this mean for simulation?

As expected, the classical approximation is at its best (approaches the quantum results) at high temperatures and for oscillators of low frequency (small *k* and/or big *m*). Under these conditions, the gaps between quantum states are small relative to thermal energy, $k_B T$.

The highest energy vibration we've looked at in detail has been that of protons oscillating at *f* = 100/ps. This frequency was chosen because it is essentially that of oxygen-hydrogen bond stretching. It represents one of the highest frequency modes of vibration in a biomolecule and thus serves as a worst-case scenario for classical approximations in macromolecular simulations. Several aspects of this motion have been depicted; see the plots of *V* and eigenvalues *E*, of probability(*x*), and the green curves in the plots of *U* and heat capacity and mean-square fluctuation.

Indeed, the mean-square fluctuation predicted classically at this frequency is about eight times too small at 300 K. We can take consolation in the scale, though. For while the quantum value of $\langle x^2 \rangle = 0.005 \text{Å}^2$ (rms *x* = 0.07 Å) is large compared to the classical result, it is still modest relative to crystallographic resolutions and the equilibrium length of the O-H bond. Furthermore, when compared to motional amplitudes measured by neutron scattering, classical simulations predict *too much* motion [6]. Thus, the reduced motion resulting from the neglect of quantum effects is overshadowed by other approximations made in simulations (perhaps the neglect of electronic polarizability and the assumed pairwise additivity of van der Waals forces). The overestimate of protein motion by simulations is not yet understood.

Another problem with classical dynamics is the incorrect partitioning of energy (see plots of *U* and heat capacity). Classically $U = k_B T$ and $C_V = k_B$, independent of frequency. In reality, high frequency motions have much more energy (larger *U*) and much less ability to exchange energy (smaller C_V) than classical mechanics predicts. High frequency motions like O-H bond stretching are energetically trapped in their quantum ground state, unexcitable to higher energy levels except at very high temperatures. Thus, the average energy *U* for the oscillator with *f* = 100/ps is very nearly temperature independent ($C_V \approx 0$) all the way up to 600 K. The frequency-dependent underestimate of *U* by classical mechanics complicates calculations of free energy differences ΔG when vibrational frequencies are likely to change during the process under investigation. However, these errors tend to cancel in estimates of $\Delta\Delta G$ values. In classical simulations, low frequency motions exchange energy with too many (high frequency) degrees of freedom, but this unphysical give-and-take of energy with high frequency motions tends to average out.

To summarize, classical simulations are unable to analyze the details of bond stretching and angle bending quantitatively. These motions are at frequencies too high for an accurate treatment using Newton's laws. However, we've observed that the errors in motional amplitude are relatively small, and errors in energy tend to cancel out in appropriately designed calculations, as when $\Delta\Delta G$'s are

calculated rather than ΔG's. For lower frequency motions ($f \approx$ 1/ps or less), observables such as U and $\langle x^2 \rangle$ become temperature independent (as quantum effects dominate) at much lower temperatures. For these motions, classical mechanics is a good approximation at physiological temperatures.

Steinbach 2003-03-07

Electrostatics and the `Generalized Born' Solvent Model

We now delve into electrostatics to estimate the *electrostatic polarization free energy*, G_{pol}, involved in the transfer of a solute with an arbitrary charge distribution from vacuum to aqueous solution. G_{pol} is the interaction between the charge distribution and its *reaction potential*, the potential induced by the charge distribution in the presence of the dielectric boundary at the solute-solvent interface.

First, we review some basics. Again, we will focus on important results, and leave most of the mathematical details to textbooks [7]. Don't worry if the equations are unfamiliar; just stay tuned for the punch line.

All problems in electrostatics boil down to the solution of a single equation, Poisson's equation:

$$\nabla^2 \phi = -\rho/\epsilon_0,$$

where ∇^2 is the *Laplacian* operator, ϕ is the *electrostatic potential*, ρ is the charge density (total charge per unit volume including all `free' and `polarization' charges), and ϵ_0 is the permittivity of free space. In cartesion (xyz) coordinates,

$$\nabla^2 \phi \equiv \frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} + \frac{\partial^2 \phi}{\partial z^2}.$$

The electrostatic potential at a given point in space is the potential energy per unit charge that a test charge would have if positioned at that point in the *electric field* $\vec{E}$ specified by ϕ:

$$\vec{E} = -\vec{\nabla}\phi \equiv -\left(\frac{\partial \phi}{\partial x}\hat{x} + \frac{\partial \phi}{\partial y}\hat{y} + \frac{\partial \phi}{\partial z}\hat{z} \right),$$

where $\hat{x}$, $\hat{y}$, and $\hat{z}$ are unit vectors in the *x*, *y*, and *z* directions, respectively. Similarly, the electric field at a given point in space is the force per unit charge that would act on a test charge located at that point. If we know ϕ at all points in space, we've solved the problem since all forces and energies can be obtained from ϕ.

Let's examine two model systems, a point charge and a point dipole, each immersed in a dielectric medium. In the following two *boundary-value* problems, we simply state the answer, giving ϕ as a function of position for all points in space. In these problems, we seek ϕ in regions were there is no charge ($\rho = 0$). Thus, we need solutions to the special case of Poisson's equation known as Laplace's equation, $\nabla^2 \phi = 0$, that satisfy two boundary conditions. First, ϕ must be a continuous function, e.g., at the dielectric boundary. Second, because there is no `free' charge (charge other than the induced polarization charges) at the dielectric boundary, the normal component of the *electric displacement*, $\vec{D} = \epsilon \vec{E}$, will also be continuous at this boundary.

First, we model a single ion in solution as a sphere of radius *a* with a point charge *q* at its center, immersed in a solvent of dielectric constant ϵ. Aside from the point charge at the center, there is nothing inside the solvent-exclusion cavity, and so the dielectric constant inside is the permittivity of free space ϵ_0. The spherical symmetry of this system renders it a problem of only one dimension, the distance *r* from the point charge. The solution is:

$$\phi(r) = \begin{cases} q/(4\pi\epsilon_0 r) - q(1/\epsilon_0 - 1/\epsilon)/(4\pi a), & r < a \\ q/(4\pi\epsilon r), & r \geq a \end{cases} \tag{6}$$

One step up in complexity from a point charge is a point dipole. So let's replace the point charge at the center of our solvent-exclusion sphere with a point dipole $\vec{p}$. With this model system we can approximate the solvation energy of a neutral molecule possessing a permanent dipole moment. Again, the dielectric constant of the solvent is ϵ, and the dielectric constant inside the spherical molecule is ϵ_0.

This cylindrically symmetric system has two independent dimensions, the distance *r* and the angle θ from the direction of the dipole vector. We get [8]:

$$\phi(r,\theta) = \begin{cases} p\cos\theta(1/r^2 - Rr)/(4\pi\epsilon_0), & r < a \\ p^*\cos\theta/(4\pi\epsilon r^2), & r \geq a \end{cases} \tag{7}$$

where

$$R = \frac{2(\epsilon - \epsilon_0)}{(2\epsilon + \epsilon_0)a^3} \quad , \quad p^* = \frac{3\epsilon}{2\epsilon + \epsilon_0}p.$$

Note that in equations 6 and 7 the potential inside the spherical molecule is a sum of two terms. In each case, the first term is the potential that would exist in the absence of the dielectric boundary at $r = a$,

and the second term is the potential induced in the spherical cavity by the charge distribution's interaction with the dielectric (e.g., the solvent). The energy of the charge distribution arising from this second term (the reaction potential) gives the electrostatic contribution to the solvation free energy.

The energy of a point charge in its reaction potential is one half of the product of the charge and the reaction potential. The `one half' appears because this is not the energy of a charge in an *external* electric field. Here, the charge has contributed to the creation of the field through its electrostatic interactions with the dielectric. So our continuum model of the solvent predicts that the electrostatic polarization free energy of solvating a spherical ion is

$$G_{ion} = \frac{1}{2}q\phi_{reaction} = -\frac{q^2}{8\pi a}\left(\frac{1}{\epsilon_0} - \frac{1}{\epsilon}\right). \tag{8}$$

This *G* is known as the Born energy [9].

The energy of a dipole in its *reaction field* (the negative gradient of the reaction potential) is minus one half of the dot product of the dipole and the reaction field. Again, this is half the energy of a dipole in an external electric field. The reaction field of the dipole is parallel to the dipole, and we get

$$G_{dipole} = -\frac{1}{2}\vec{p}\cdot\vec{E}_{reaction} = -\frac{p^2}{4\pi\epsilon_0 a^3}\frac{(\epsilon - \epsilon_0)}{(2\epsilon + \epsilon_0)}. \tag{9}$$

Several names (Bell, Onsager, Kirkwood) have been associated with this energy.

Note that both *G* values are zero if $\epsilon = \epsilon_0$, i.e., if we haven't changed the dielectric constant of the environment.

Still and coworkers [10] have proposed the following approximate expression for the free energy of solvent polarization for an arbitrary charge distribution of *N* charges:

$$G_{pol} = -\frac{1}{8\pi}\left(\frac{1}{\epsilon_0} - \frac{1}{\epsilon}\right)\sum_{i,j=1}^{N}\frac{q_i q_j}{f_{GB}}, \tag{10}$$

where

$$f_{GB} = \sqrt{r_{ij}^2 + a_{ij}^2 e^{-D}}, \quad D = r_{ij}^2/(2a_{ij})^2, \quad a_{ij} = \sqrt{a_i a_j}.$$

This functional form of the so-called Generalized Born (GB) approximation has been used with considerable success to efficiently evaluate hydration energies for small molecules. Parameterization of the method involves accounting for the effects of neighboring solute atoms in the determination of each atom's effective Born radius *a*.

As shown in the following figure, this GB approximation behaves appropriately in important limiting situations. For *N* identical, coincident ($r_{ij} = 0$) particles of charge *q*, it gives the correct Born energy

(equation $\underline{8}$, for a single particle of charge Nq). For two charges of equal and opposite sign, it approaches the dipole result (equation $\underline{9}$) at short separation distances, as it should. For two well separated charges ($r_{ij} > 2.5a_{ij}$), it approaches the appropriate energy: the two Born energies plus the energetic change in the Coulomb interaction between the two charges due to the dielectric medium.

GB approximation (in red) to solvent polarization energies for two charges of equal radii as a function of separation. Upper curves: Equal and opposite charges with *G_dipole* (blue) at small separation and Coulomb + Born polarization energies (green) at large separation. Lower curves: Equal charges with *G_ion* (blue) at zero separation and Coulomb + Born polarization energies (green) at large separation.

 ı   previc

Steinbach 2003-03-07

Classical Macromolecular Simulation

To simulate the structure and dynamics of biomolecules, we approximate them as *a physical network of balls that have point charges at their centers and are connected by springs*. In addition to springs that govern the bending of bonds and angles, there are forces that favor certain rotations about the bonds. The balls representing the atoms are not hard spheres; they are Lennard-Jones particles that can overlap each other. Our goal is to study the motion of this physical network of balls and springs, in hopes of interpreting and predicting the dynamics of real macromolecules at the atomic level.

Subsections

- A Note on Notation - Some CHARMMing language
- Simulating Biomolecules: The Three Necessary Ingredients
    - I. A Description of the Structure
    - II. Initial Coordinates
    - III. An Empirical Energy Function: Free Energy vs. Potential Energy

---

Steinbach 2003-03-07

A Note on Notation - Some CHARMMing language

The following discussion of classical simulation techniques is quite general. But when introducing a particular concept, its name in the program CHARMM may be given parenthetically. Every program designed to simulate macromolecules deals with these concepts in its own way. Connecting the concepts to CHARMM-like nomenclature simply reflects my experience and may better clarify the concepts for those with particular interest in this simulation package.

Steinbach 2003-03-07

Simulating Biomolecules: The Three Necessary Ingredients

To simulate the dynamics of a macromolecule, we need to specify three things:

Subsections

Steinbach 2003-03-07

I. A Description of the Structure

The **Protein Structure File (CHARMM: `.psf')** contains a description of all molecules in the system to be simulated, their covalent connectivity, and all the energetic interactions to be calculated. Despite its name, the psf is not limited to proteins. It may be used to describe other systems, such as nucleic acids or lipids. The psf defining a *specific* system (e.g., carboxymyoglobin (MbCO) hydrated by 350 water molecules) is assembled, or GENErated, from the *generic* building blocks defined in the **Residue Topology File (.rtf)**. To generate the psf for hydrated MbCO, we need amino acids and water molecules (defined in aminoh.rtf), and a heme and CO ligand (in porphyrinh.rtf). Again, despite the term `rtf', the components found in an rtf need not be amino acid residues. Other rtf's contain different building blocks: base pairs, lipid head groups, etc.

Steinbach 2003-03-07

II. Initial Coordinates

The psf only specifies what atoms are involved, how they're connected, and what interactions determine the system's energy. Next, we need to place the atoms somewhere. Simulations are generally begun with experimentally determined coordinates, typically from the Brookhaven Protein Data Bank. Structures determined by X-ray diffraction lack positions for hydrogen atoms, while those determined by neutron diffraction (less common) or nuclear magnetic resonance do not. For X-ray structures, the hydrogen positions can be assigned by CHARMM (using the HBUIld command). Initial coordinates of small systems can also be created in the absence of experimental data by invoking ideal stereochemistry, as specified in the parameter file. This can also be done for proteins for which limited data exist, e.g., if only the C_α coordinates are known to within 1 Å or so.

Steinbach 2003-03-07

III. An Empirical Energy Function: Free Energy vs. Potential Energy

Finally, given a defined system and its initial atomic coordinates, we need a function describing the energy of the system for any configuration, $\vec{R}$, of the atomic coordinates. A functional form must be chosen for the energy as well as the associated numerical constants. For macromolecular simulation potentials, these **parameters** number in the thousands and include spring stiffnesses and equilibrium distances, torsional barriers and periodicities, partial charges, and Lennard-Jones coefficients. The energy function and its associated constants are contained in the **Parameter File (.prm)**. Development of parameter sets is a laborious process. Both the functional form and numerical parameters require extensive optimization. A brute-force iteration of simulation and parameter modification is performed to improve agreement between simulations of model systems and information derived from *ab initio* calculations, small-molecule spectroscopy, and educated guessing.

Free Energy vs. Potential Energy

For a system held at constant NVT, the **Helmholtz free energy**, $A \equiv U - TS = -kT\ln(Z)$, is a minimum at equilibrium, where $U = \langle E \rangle$ is the average total energy of the system (kinetic energy plus *potential energy*), T is the absolute temperature, S is the entropy, and Z is the partition function (eq 3). Suppose the pressure P is held constant instead of the volume (constant NPT, the `isothermal isobaric' ensemble). In this case, the **Gibbs free energy**, $G \equiv U + PV - TS$, is minimized at equilibrium. Note that the *enthalpy*, $H = U + PV$, is the quantity at constant pressure that corresponds to U at constant volume. Differences in G drive chemical reactions.

Some empirical energy functions are designed to approximate the Gibbs free energy G. For example, in Monte Carlo studies of protein structure prediction, the energy function may be based simply on the likelihood of residues of type i and j being within a certain distance of each other. The probabilities p are determined by counting the number of times that residues i and j are found close to each other in the protein structures deposited in the Protein Data Bank. They are then converted into ΔG-like energies by: $p_{ij} \propto e^{-\Delta G_{ij}/k_B T}$. Because the p's are derived from structures at constant T and P determined experimentally, these energy functions account for entropic contributions to the Gibbs free energy in an approximate way.

In most molecular dynamics software packages, however, the empirical energy function, $V(\vec{R})$ (not to be confused with the volume V), is developed to approximate the *potential energy* of the system. In general, it does not include entropic effects in any effective way. Many simulations have been performed at constant energy, E. That is, E is fixed and T fluctuates about an average value as energy is exchanged between the kinetic energy and the potential energy. In principle, simulations performed at constant T and P mimic experimental conditions better than simulations at constant E. Recently, an improved constant-PT algorithm has been developed [11]. Constant E simulations have the advantage

that they allow energy conservation to be checked. Any significant drifts in E indicate a problem that should be tracked down before continuing the simulation. Although it fluctuates, the temperature is still well defined at constant E, and differences between dynamics at constant T and constant E are generally not too significant on the time scales currently accessible to MD simulation (100's of ps to a few ns). However, the constant-PT simulation may well become the standard as large solvated systems are simulated over longer time scales.

Steinbach 2003-03-07

The Empirical Potential Energy Function

Each of the interactions commonly employed in the potential energy function $V(\vec{R})$ is sketched below.

Simple harmonic terms describe **bond stretching** and **angle bending**. The planarity of groups (e.g., the amide planes of proteins) can also be enforced by harmonic potentials known as an **improper dihedrals**. Rotation about single bonds (**torsions**) is governed by sinusoidal energies.

The **electrostatic attraction or repulsion** between two charges is described by Coulomb's law:

$$V_{ij}^{Coulomb}(r_{ij}) = \frac{q_i q_j}{4\pi\epsilon_r\epsilon_0 r_{ij}},$$

where q_i and q_j are the atoms' partial charges, r_{ij} is the distance separating the atoms' centers, ϵ_0 is the permittivity of free space, and ϵ_r is the relative dielectric coefficient of the medium between the charges (i.e., $\epsilon_{medium} = \epsilon_r\epsilon_0$).

A distance-dependent dielectric coefficient (RDIE: $\epsilon_r = r_{ij}$) has been used to approximate solvent screening without including explicit water molecules. Physically, it's a pretty ugly way to cheat. But if you don't want to include water, it may be the best your simulation package has to offer; it is almost certainly better than using unscreened partial charges in the absence of water. For realistic dynamics, we recommend constant-dielectric (CDIE) simulations with explicit solvation and $\epsilon_r = 1$. The presence of water retards conformational searching, however.

An important electro*dynamic* effect remains to be included: **van der Waals interactions**. The electron cloud of a neutral atom fluctuates about the positively charged nucleus. The fluctuations in neighboring atoms become correlated, inducing attractive dipole-dipole interactions. The equilibrium distance between two proximal atomic centers is determined by a trade off between this attractive dispersion force and a core-repulsion force that reflects electrostatic repulsion and the Pauli exclusion principle. The Lennard-Jones potential models the attractive interaction as $\propto r_{ij}^{-6}$ and the repulsive one as $\propto r_{ij}^{-12}$:

$$V_{ij}^{LJ}(r_{ij}) = \varepsilon_{ij}\left[\left(\frac{r_{ij}^{min}}{r_{ij}}\right)^{12} - 2\left(\frac{r_{ij}^{min}}{r_{ij}}\right)^6\right],$$

where r_{ij}^{min} is the equilibrium separation distance (where the force $F = -dV_{ij}^{LJ}/dr_{ij} = 0$) and ε_{ij} is the well depth; i.e., $V_{ij}^{LJ}(r_{ij}^{min}) = -\varepsilon_{ij}$. Why this `6-12' form for the van der Waals interaction? The application of quantum perturbation theory to two well separated hydrogen atoms in their ground states yields an interaction energy that decays as r_{ij}^{-6}, and r_{ij}^{-12} is obviously easy to calculate from

r_{ij}^{-6}. For simplicity, the Lennard-Jones forces are typically modeled as effectively pair-wise additive: the potential energy V_{ABC} of three adjacent particles A, B, and C is the sum of the three energies for each atom pair: $V_{ABC} = V_{AB} + V_{BC} + V_{AC}$. Pair-wise additivity is only an approximation.

Perhaps, you are thinking, `*Hey, what about magnetic forces?*' The magnetic force between two moving charges is expressed in terms of a double vector cross product involving the two particle velocities and the vector $\vec{r}_{ij}$ of separation. It does not generally act along $\vec{r}_{ij}$, but it does when two charges q have instantaneous velocities v along parallel lines. For this case, we can conveniently compare the magnitudes of the magnetic and electric forces. It turns out that the magnetic force is weaker than the electric force by a factor of $(v/c)^2$, where c is the speed of light. Thus, magnetic forces are neglible for nonrelativistic particles, such as the partial charges that are used in simulation force fields. For example, if a particle moves as much as 1 Å in as short a time as 1 femtosecond (10^{-15} s), then $(v/c)^2 = 1.1 * 10^{-7}$. We may therefore completely neglect magnetic interactions.

# Empirical Potential Energy Function

Bonds

Angles

Improper Dihedrals

Torsions

Electrostatics

van der Waals

**Interactions included in representative potential energy function
for MD simulation.**

For those who like equations with their pictures, a typical potential energy function used in MD simulations looks like:

$$V(\vec{R}) = V_{bonded}(\vec{R}) + V_{nonbonded}(\vec{R}),$$

with

$$V_{bonded}(\vec{R}) = \sum_{bonds} k_l(l-l_0)^2 + \sum_{angles} k_\theta(\theta-\theta_0)^2 + \sum_{impropers} k_\omega(\omega-\omega_0)^2 + \sum_{torsions} A_n\left[1+cos(n\phi\right.$$

$$V_{nonbonded}(\vec{R}) = \sum_{i<j}\left(\varepsilon_{ij}\left[\left(\frac{r_{ij}^{min}}{r_{ij}}\right)^{12} - 2\left(\frac{r_{ij}^{min}}{r_{ij}}\right)^6\right] + \frac{q_iq_j}{4\pi\epsilon_r\epsilon_0 r_{ij}}\right)$$

The first `bonded' sum is over bonds between atom pairs; the second sum is over bond angles defined by three atoms; the third and fourth sums are over atom foursomes (as in the figure above). For bookkeeping purposes, each atom is assigned a number. In the `nonbonded' interactions (van der Waals and electrostatics), the summation is over atoms i and j, where `$i < j$' simply ensures that each interaction is counted only once. Generally, atoms separated by one or two bonds are excluded from the nonbonded sum, and those separated by three bonds, `1-4 interactions', may have electrostatic interactions reduced by a multiplicative scale factor. The form of $V(\vec{R})$ shown here reflects the choice not to include an

explicit hydrogen bond term, favoring instead to account for hydrogen bonds through an appropriate parameterization of Lennard-Jones and Coulomb interactions. Note also that a single dihedral angle (torsion) may have an energy described by more than one Fourier component (multiple values of n).

Subsections

- Matching CHARMM's Electrostatic Approximations to Environmental Approximations
  - No Explicit Solvent (RDIE)
  - Partial Solvation by Explicit Water Molecules
  - Solutions, Crystals, and Interfaces


Steinbach 2003-03-07

Matching CHARMM's Electrostatic Approximations to Environmental Approximations

To speed up computation of $V(\vec{R})$, electrostatic and van der Waals forces are commonly terminated at

a specified distance (CTOFnb) from the atom exerting the forces [12]. For van der Waals forces, use the CHARMM keyword VSHIft. For electrostatic forces, the choice of approximation depends on the system to be simulated. Ewald summation is being used more and more as a way to treat electrostatics without any cutoff in periodically repeated (infinite) systems such as solutions or crystals. However, the appropriate choice of spherical cutoff is still relevant for any modeling and simulation of a finite system, e.g., a hydrated protein in vacuum.

The following figure shows the potential energy of electrostatic interaction for two unit charges (top), as approximated by various methods using a cutoff distance of 12 Å. Also shown is the error in the Coulomb force (bottom) resulting from these approximations. Note the large force errors at long range obtained when using a switching function on the potential energy (*switch*). **This potential-switching method should not be used.** The choice made from among the other alternatives available in CHARMM should be made based on the system simulated, as detailed in the following discussion.

Potential energy and error in force calculated by some `spherical cutoffs.'

Subsections

- [No Explicit Solvent (RDIE)](#)
- [Partial Solvation by Explicit Water Molecules](#)
- [Solutions, Crystals, and Interfaces](#)

Steinbach 2003-03-07

No Explicit Solvent (RDIE)

A distance-dependent dielectric coefficient (RDIE: $\epsilon_r = r_{ij}$) has been used to approximate solvent

screening without including explicit water molecules. Physically, it's a pretty ugly way to cheat. But if you don't want to include water, it may be the best your simulation package has to offer; it is almost certainly better than using unscreened partial charges in the absence of water. For realistic dynamics, we recommend constant-dielectric (CDIE) simulations with explicit solvation and $\epsilon_r = 1$. The presence of

water retards conformational searching, however.

The SHIFt option does a good job of monotonically damping a $1/r^3$ force (RDIE electrostatics) to

zero. Use of CTOFnb = 10 Åshould suffice. Although simulations have definitely evolved toward explicit solvation with true Coulombic $1/r^2$ forces (CDIE), RDIE SHIFt is arguably the best choice (in

CHARMM) when the inclusion of explicit water remains prohibitively expensive.

Steinbach 2003-03-07

Partial Solvation by Explicit Water Molecules

When water molecules are included in the simulation, constant-dielectric (CDIE) electrostatics should be employed. For systems of finite size (e.g., a protein hydrated in vacuum), spherical cutoffs are used. Good numerical input values are CTOFnb ≥ 12 (Å), CUTNb $=$ CTOFnb $+$ 2, and INBFreq $= -1$.

Electrostatic SHIFting does not monotonically damp the $1/r^2$ force, but force shifting (FSHIft) and

force switching (FSWItch) do. If the system simulated is composed entirely of neutral groups, force shifting (ATOM FSHIft) is the better of the two. As the number of charged groups in the system increases (e.g., hydrated protein), force switching (ATOM FSWItch) is recommended with CTONnb $=$ CTOFnb $-$ 4.

Steinbach 2003-03-07

## Solutions, Crystals, and Interfaces

Solutions and crystals are simulated as `infinite' systems. More precisely, the basic `unit cell' (e.g., a single protein in a cube of water) having zero net charge is replicated periodically in all three dimensions. Periodic boundary conditions are used so that an atom exiting through one face of the unit cell enters the cell through the opposite face. These simulations are expensive, requiring considerable explicit water and the calculation of forces exerted by the `image' atoms that neighbor the primary unit cell. Ewald summation is the preferred electrostatic treatment for these periodic systems. If an interface (e.g., air-water) is present, Ewald summation may improve results even more than otherwise.

Steinbach 2003-03-07

Questions We Can Ask With a Computer

Given a structural description and atomic coordinates: `What does it look like?'

This is an obvious question to ask, but the value of sitting down and staring at a structure is difficult to overestimate. Of course, the pretty pictures generated by graphics packages are only representations of models, even the pretty pictures of `experimental' structures. CHARMM graphics has broad capabilities but is less convenient than some commercial packages, such as QUANTA or SYBYL. It's hard to beat RASMOL, a public domain program that spins proteins on command, with options to display them as ribbons, balls and sticks, or space-filling overlapping spheres. RASMOL also allows point-and-click atom identification as well as limited zooming and z-clipping.

Given a structural description, atomic coordinates, and an energy function: `How does the system relax and fluctuate?'

Now we're getting to the point! Structure determination is clearly a critical step toward understanding biological function, but protein function requires motion. Molecular dynamics is the link between structure and function.

We might, for example, wish to characterize the dependence of a protein's structure and dynamics on environmental conditions. We could perform simulations at different temperatures, different pressures, or different levels of hydration. We could approximate the solution environment by a periodically repeating system in which the repeating unit was a single protein in a box of water. Or the crystalline phase could be simulated as a special case of the periodic system with a particular box size and shape.

To answer questions like these on a computer, we need to employ a few techniques that manipulate the structure, $\vec{R}$, given the potential energy, $V(\vec{R})$.

Steinbach 2003-03-07

Energy Minimization

Function optimization is a calculation that pervades much of numerical analysis. In the context of macromolecules, the function to be optimized (minimized) is an energy. The energy landscape of a biomolecule possesses an enormous number of minima, or conformational substates. Nonetheless, the goal of energy minimization is simply to find the the local energy minimum, i.e., the bottom of the energy well occupied by the initial conformation (★ in figure). The energy at this local minimum may be much higher than the energy of the global minimum. Physically, energy minimization corresponds to an instantaneous freezing of the system; a static structure in which no atom feels a net force corresponds to a temperature of 0 K. In the early 1980's, energy minimization was about all one could afford to do and was dubbed `molecular mechanics.'
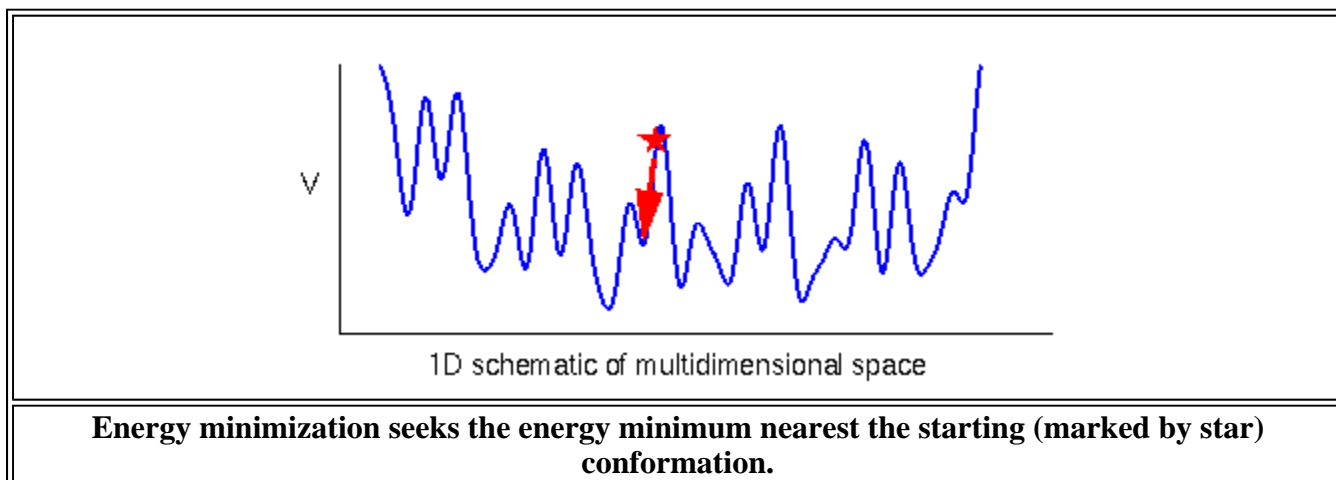
Energy minimization seeks the energy minimum nearest the starting (marked by star) conformation.

Steinbach 2003-03-07

Molecular Dynamics (MD) Simulation

As already noted, MD simulation generally begins where experimental structure determination leaves off, if not during the structure refinement itself. It is generally not used to predict structure from sequence or to model the protein folding pathway. MD simulation can fold extended sequences to `global' potential energy minima for very small systems (peptides of length ten, or so, in vacuum), but it is most commonly used to simulate the dynamics of known structures.

An initial velocity is assigned to each atom, and Newton's laws are applied at the atomic level to propagate the system's motion through time (see `Classical and Quantum Mechanics - in a Nutshell' above). Thus, dynamical properties such as time correlation functions and transport coefficients (e.g., diffusion constants, bulk viscosities) can be calculated from a sufficiently long MD trajectory.

Once again, Newton's second law is: $\vec{F}_i = m_i \vec{a}_i$, where $\vec{F}_i$ is the sum of all forces acting on atom i

that results in its acceleration $\vec{a}_i$. The acceleration is the second derivative of the position with respect

to time: $\vec{a}_i = d\vec{v}_i/dt = d^2\vec{r}_i/dt^2$. In words, it is the rate of change of the velocity $\vec{v}_i$, which in turn,

is the rate of change of the position $\vec{r}_i$.

The `Leap Frog' algorithm is one method commonly used to numerically integrate Newton's second law. We obtain all atomic positions $\vec{r}_i$ at all times t_n and all atomic velocities $\vec{v}_i$ at intermediate times

$t_{n+1/2}$. This method gets its name from the way in which positions and velocities are calculated in an alternating sequence, `leaping' past each other in time:

$$\vec{v}_i(t_{n+1/2}) = \vec{v}_i(t_{n-1/2}) + \frac{\vec{F}_i(t_n)}{m_i}\Delta t,$$

$$\vec{r}_i(t_{n+1}) = \vec{r}_i(t_n) + \vec{v}_i(t_{n+1/2})\Delta t.$$

Initial velocities are assigned so as to reflect equilibrium at the desired temperature T (a Maxwellian distribution), without introducing a net translation or rotation of the system.

The energy of an isolated system (as opposed to, for example, one in contact with a thermal bath) is conserved in nature, but it may not be in simulations. Energy conservation can be violated in simulations because of an insufficiently short integration time step Δt, an inadequate cutoff method applied to

long-range (electrostatic and Lennard-Jones) forces, or even bugs in the program. Of course, energy conservation alone is not sufficient to ensure a realistic simulation. The realism of the dynamics trajectory depends on the empirical potential energy function $V(\vec{R})$, the treatment of long-range forces, the value of Δt, etc.

Steinbach 2003-03-07

Langevin Dynamics (LD) Simulation

The Langevin equation is a *stochastic* differential equation in which two force terms have been added to Newton's second law to approximate the effects of neglected degrees of freedom. One term represents a frictional force, the other a *random* force $\vec{R}$. For example, the effects of solvent molecules not explicitly present in the system being simulated would be approximated in terms of a frictional drag on the solute as well as random kicks associated with the thermal motions of the solvent molecules. Since friction opposes motion, the first additional force is proportional to the particle's velocity and oppositely directed. Langevin's equation for the motion of atom *i* is:

$$\vec{F}_i - \gamma_i \vec{v}_i + \vec{R}_i(t) = m_i \vec{a}_i,$$

where $\vec{F}_i$ is still the sum of all forces exerted on atom *i* by other atoms explicitly present in the system.

This equation is often expressed in terms of the `collision frequency' $\zeta = \gamma/m$.

The friction coefficient is related to the fluctuations of the random force by the *fluctuation-dissipation theorem*:

$$\langle \vec{R}_i(t) \rangle = 0,$$

$$\int \langle \vec{R}_i(0) \cdot \vec{R}_i(t) \rangle dt = 6 k_B T \gamma_i.$$

In simulations it is often assumed that the random force is completely uncorrelated at different times. That is, the above equation takes the form:

$$\langle \vec{R}_i(t) \cdot \vec{R}_i(t') \rangle = 6 k_B T \gamma_i \delta(t - t').$$

The temperature of the system being simulated is maintained via this relationship between $\vec{R}(t)$ and γ.

The jostling of a solute by solvent can expedite barrier crossing, and hence Langevin dynamics can search conformations better than Newtonian molecular dynamics ($\gamma = 0$).
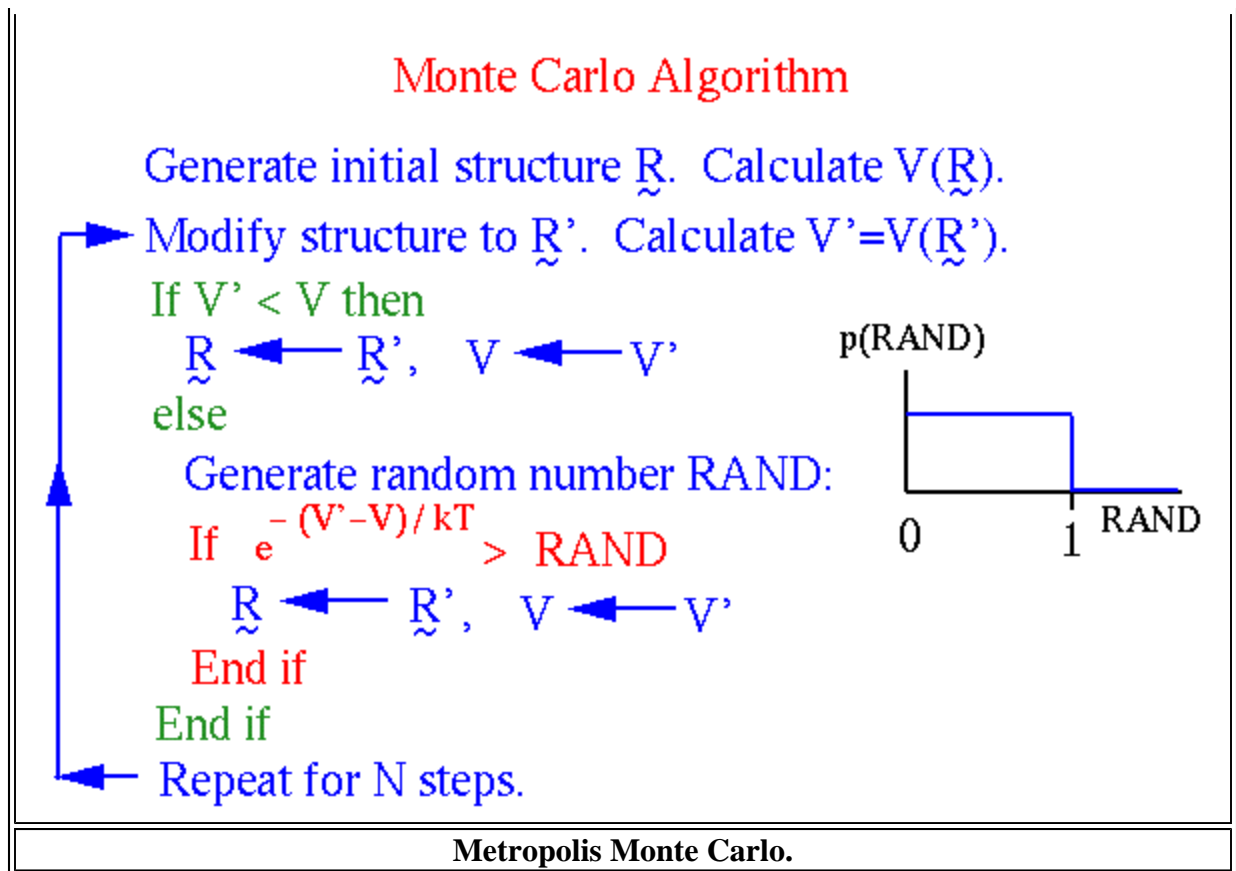
Monte Carlo (MC) Simulation

Instead of evaluating forces to determine incremental atomic motions, Monte Carlo simulation simply imposes relatively large motions on the system and determines whether or not the altered structure is energetically feasible at the temperature simulated. The system jumps abruptly from conformation to conformation, rather than evolving smoothly through time. It can traverse barriers without feeling them; all that matters is the relative energy of the conformations before and after the jump. Because MC simulation samples conformation space without a true `time' variable or a realistic dynamics trajectory, it cannot provide time-dependent quantities. However, it may be much better than MD in estimating average thermodynamic properties for which the sampling of many system configurations is important.


Monte Carlo makes use of Boltzmann probabilities, not forces.
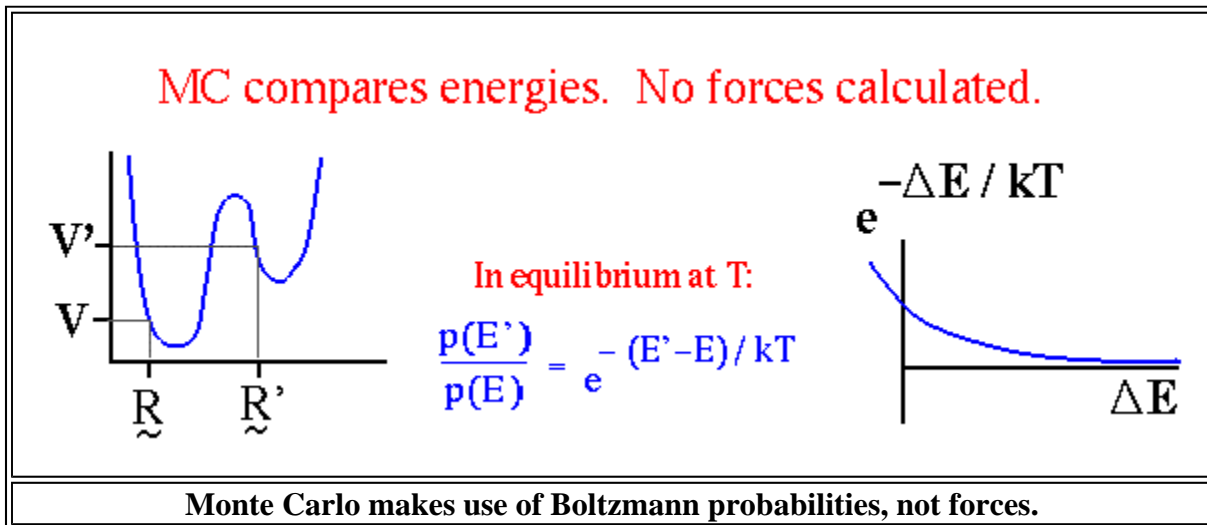
When the potential energy *V* and observables to be calculated from the simulation are velocity-independent (as is typical), an MC simulation need only compare potential energies *V*, not total energies *E* (see `Calculating Equilibrium Averages' above). Two conformations, $\vec{R}$ and $\vec{R'}$, are compared and updated as shown below [13]. $RAND$ is a random number uniformly distributed on [0,1].

Monte Carlo Algorithm

Generate initial structure $\underset{\sim}{R}$. Calculate $V(\underset{\sim}{R})$.

Modify structure to $\underset{\sim}{R}'$. Calculate $V' = V(\underset{\sim}{R}')$.

If $V' < V$ then

$\underset{\sim}{R} \longleftarrow \underset{\sim}{R}'$, $V \longleftarrow V'$

else

Generate random number RAND:

If $e^{-(V'-V)/kT} > $ RAND

$\underset{\sim}{R} \longleftarrow \underset{\sim}{R}'$, $V \longleftarrow V'$

End if

End if

Repeat for N steps.

Metropolis Monte Carlo.

For simple systems, the structural modifications are often tuned so that about 50% of the $\vec{R}'$ conformations are accepted. For macromolecular systems, this acceptance ratio can be much smaller, e.g. when dihedral angles are modified by large amounts. It is then generally expedient to bias the random moves in favor of known structural preferences such as side chain rotamers (`biased probability Monte Carlo') and to do some energy minimization before evaluating the energy V'. Also, explicit water molecules hinder the acceptance of new conformations. So Monte Carlo simulations of macromolecules generally use an implicit model of solvation, e.g., a term in the empirical potential energy function that mimics the effects of water. See, for example, reference [14].

Steinbach 2003-03-07

Normal Mode (Harmonic) Analysis

Normal modes of vibration are simple harmonic oscillations about a *local* energy minimum, characteristic of a system's structure $\vec{R}$ and its energy function $V(\vec{R})$. For a purely harmonic $V(\vec{R})$, any motion can be exactly expressed as a superposition of normal modes. For an anharmonic $V(\vec{R})$, the potential near the minimum will still be well approximated by a harmonic potential, and any small-amplitude motion can still be well described by a sum of normal modes. In other words, at sufficiently low temperatures, any classical system behaves harmonically.

In a typical normal mode analysis, the characteristic vibrations of an energy-minimized system ($T = 0^+$ K) and the corresponding frequencies are determined assuming $V(\vec{R})$ is harmonic in all degrees of freedom. Normal mode analysis is less expensive than MD simulation, but requires much more memory.

As a globular protein is heated from very low temperature, the fluctuations of its atoms begin to deviate measurably from harmonic behavior around 200 K. The motion at 300 K is considerably anharmonic. This must be kept in mind when attempting to interpret physiological behavior in terms of normal modes. Still, calculation of the normal mode spectrum is less expensive than a typical MD simulation, and the spectrum may provide qualitative, if not quantitative, insight.

The normal mode spectrum of a 3-dimensional system of N atoms contains $3N - 6$ normal modes ($3N - 5$ for linear molecules in 3D). In general, the number of modes is the system's total number of degrees of freedom minus the number of degrees of freedom that correspond to pure rigid body motion (rotation or translation). Each mode is defined by an *eigenvector* and its corresponding *eigenfrequency, ω*. The eigenvector contains the amplitude and direction of motion for each atom. In mode *i*, all N atoms oscillate at the same frequency, ω_i.

Examples of Normal Modes

Guitar string fixed at both ends

$$\lambda = 2L / n$$
$$n = 1, 2, 3$$

Balls & Springs in 1D

| ω	mode
0	(Pure Translation)
$(k/m)^{1/2}$	(Center at Rest)
$[(k/m)(1+2m/M)]^{1/2}$	

Proteins

Low frequency GLOBAL

High frequency LOCAL

Small oscillations about an equilibrium position can be expressed in terms of normal modes.

In macromolecules, the lowest frequency modes correspond to delocalized motions, in which a large number of atoms oscillate with considerable amplitude. The highest frequency motions are more localized, with appreciable amplitudes for fewer atoms, e.g., the stretching of bonds between carbon and hydrogen atoms.
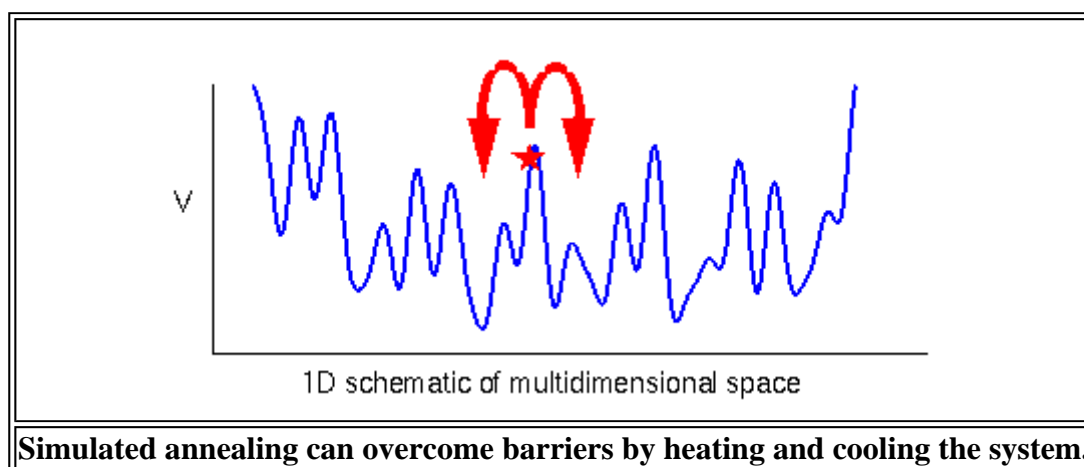
Simulated Annealing

Simulated annealing is a special case of either MD (`quenched' MD), LD, or MC simulation, in which the temperature is gradually reduced during the simulation. Often, the system is first heated and then cooled. Thus, the system is given the opportunity to surmount energetic barriers in a search for conformations with energies lower than the local-minimum energy found by energy minimization. This improved equilibration can lead to more realistic simulations of dynamics at low temperature [6]. Of course, annealing is more expensive than energy minimization. Simulated annealing is often applied to

potentials, $V(\vec{R})$, that include unphysical energy terms, as when annealing structures to reduce

crystallographic R factors.

Simulated annealing can overcome barriers by heating and cooling the system.

Steinbach 2003-03-07

What is Unique to Computer Experiments?

- May look very closely ($\vec{r}_i(t)$, $\vec{v}_i(t)$) at the behavior of *any atomic subset* of the system. In

  principle, any function of the atomic positions and velocities, whether time-averaged or instantaneous, is computable.

- May modify potential function $V(\vec{R})$ arbitrarily. Examination of what nature does not do may

  provide a new understanding for what it does do. For example, we have investigated the contribution of torsional transitions to the anharmonicity of protein dynamics by comparing simulations of MbCO dynamics performed with and without infinitely high barriers that prohibit these transitions. Our conclusion: Dihedral transitions account for nearly all the motional anharmonicity of dried MbCO but for less than half of the motional anharmonicity of hydrated MbCO [16].

- May mutate structures or environments slowly (Free Energy Perturbation Theory) and approximate differences in free-energy differences, $\Delta\Delta G$. By taking differences in ΔG, the states of zero energy are consistently defined and errors due to approximations in the simulation protocol tend to subtract out. See `Classical vs. Quantum Mechanics: The Harmonic Oscillator in One Dimension' above.

Steinbach 2003-03-07

Worth Worrying About:

Ultimately, simulations are judged according to two basic criteria:

I. How well do the empirical energy surface and the chosen system composition approximate Nature?

$V(\vec{R})$:

> How realistic are the chosen functional form and the associated numerical constants?

PSF Generation:

> Which titratable groups should be protonated? Without employing quantum mechanics, protonations are assumed at the beginning and maintained throughout the simulation. Also, how much water is needed [15]? How many ions should be included?

II. How well is the energy surface (phase space) explored?

MD Simulation:

> What length of simulation is sufficient? First, the system must be equilibrated such that system properties such as potential energy, temperature, and volume appear to have stopped drifting. Then the simulation must continue long enough to obtain reliable equilibrium averages.

MC Simulation:

> Does the chosen `move set' embody all motions relevant to the question being asked of the simulation? Have enough steps been taken?

<div align="center">Mistakes to Avoid:</div>

Inconsistent $V(\vec{R})$:

> The potential function (long-range cutoff keywords, distances, ...) should not be changed at different stages of a simulation study. All input scripts used in a research project that evaluate energies and forces (energy minimizations, annealings, dynamics simulations, ...) should *explicitly* (Don't trust the defaults!) do so in the same way.

Submit and Forget:

> Don't let a simulation run unmonitored. Check intermediate results daily. Plot the time dependences of the potential and total energies, the temperature, the pressure and volume (if applicable), and the root-mean-square deviation from a reference (crystallographic or $t = 0$) structure.

<div align="center">Remember:</div>

Simulations are fiction aspiring to emulate reality. Pretty pictures and even a few good numbers do not guarantee good science.

**Next:** Bibliography **Up:** intro_simulation **Previous:** What is Unique to *Steinbach 2003-03-07*

Bibliography

1   L.I. Schiff, `Quantum Mechanics', 3rd ed., (McGraw-Hill, New York, 1968), chap 1,2.

2   L. Pauling and E.B. Wilson, Jr., `Introduction to Quantum Mechanics', (McGraw-Hill, New York, 1935), chap 3.

3   W. Doster, S. Cusack and W. Petry, (1989) *Nature* **337**, 754-756.

4   Krupyanskii, Y.F., Parak, F., Goldanskii, V.I., Mössbauer, R.L., Gaubman, E.E., Engelmann, H., & Suzdalev, I.P. (1982) *Z. Naturforsch., C: Biosci.* **37C**, 57-62.

5   B.R. Brooks, R.E. Bruccoleri, B.D. Olafson, D.J. States, S. Swaminathan and M. Karplus, (1983) *J. Comp. Chem.* **4**, 187-217.

6   P.J. Steinbach and B.R. Brooks, (1994) *Chem. Phys. Letters* **226**, 447-452.

7   R.P. Feynman, R.B. Leigthton and M. Sands, `The Feynman Lectures on Physics', Vol. II, (Addison-Wesley, Reading, 1965), chap 4-10.

8   L. Onsager, (1936) *J. Am. Chem. Soc.* **58**, 1486-1493.

9   M. Born, (1920) *Z. Phys.* **1**, 45.

10   W.C. Still, A. Tempczyk, R.C. Hawley and T. Hendrickson, (1990) *J. Am. Chem. Soc.* **112**, 6127-6129.

11   S.E. Feller, Y. Zhang, R.W. Pastor and B.R. Brooks, (1995) *J. Chem. Phys.* **103**, 4613-4621.

12   P.J. Steinbach and B.R. Brooks, (1994) *J. Comp. Chem.* **15**, 667-683.

13   N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller and E. Teller, (1953) *J. Chem. Phys.* **21**, 1087-1092.

14   R. Abagyan and M. Totrov, (1994) *J. Mol. Biol.* **235**, 983-1002.

15   P.J. Steinbach and B.R. Brooks, (1993) *Proc. Natl. Acad. Sci. USA* **90**, 9135-9139.

16   P.J. Steinbach and B.R. Brooks, (1996) *Proc. Natl. Acad. Sci. USA* **93**, 55-59.

Steinbach 2003-03-07

   previc

Up: intro_simulation **Previous:** Bibliography

About this document ...

This document was generated using the **LaTeX**2HTML translator Version 2K.1beta (1.61)

Copyright © 1993, 1994, 1995, 1996, Nikos Drakos, Computer Based Learning Unit, University of Leeds.
Copyright © 1997, 1998, 1999, Ross Moore, Mathematics Department, Macquarie University, Sydney.

The command line arguments were:
latex2html -image_type gif intro_simulation

The translation was initiated by Steinbach on 2003-03-07

Steinbach 2003-03-07