

Resolution exchange simulation with incremental coarsening

Edward Lyman*and Daniel M. Zuckerman†
Department of Computational Biology, School of Medicine
and Department of Environmental and Occupational Health,
Graduate School of Public Health,
BST W1041, 200 Lothrop St., University of Pittsburgh, Pittsburgh, PA 15261

March 9, 2006

Abstract

We previously developed an algorithm, called “resolution exchange”, which improves canonical sampling of atomic resolution models by swapping conformations between high- and low-resolution simulations.¹ Here, we demonstrate a generally applicable incremental coarsening procedure and apply the algorithm to a larger peptide, met-enkephalin. In addition, we demonstrate a combination of resolution and temperature exchange, in which the coarser simulations are also at elevated temperatures. Both simulations are implemented in a “top-down” mode, to allow efficient allocation of CPU time among the different replicas.

Atomic resolution simulations of proteins are currently limited to short durations (less than one μsec)² or small systems (less than 100 residues).^{3,4} Furthermore, accurate calculations involving large conformational changes are not possible for any system, as the cost of calculating entropic contributions is too great. Indeed, the cost of such calculations is only going to increase, as empirical force fields are improved by including polarization effects, either in a classical way^{5,6} or in a semiclassical way.⁷

Thoroughly sampling the space of conformations is essential for a number of problems. From a purely biological perspective, there is a growing awareness of the importance of protein fluctuations—over and above

*elyman@ccbb.pitt.edu

†dmz@ccbb.pitt.edu

the static picture—in the function of most proteins.⁸ Allostery and conformational changes dramatic enough to be captured experimentally are just two examples of the existence of such fluctuations.^{9,10} In a computational context, careful validation of empirical forcefields requires confidence in the quality of conformational sampling, so that error may be attributed to the forcefield rather than undersampling. The calculation of free energy differences—as required for evaluation of binding affinities of small molecules,¹¹ or the strength of protein-protein interactions¹²—also requires reliable sampling.^{13,14}

The undersampling (or “quasi-ergodicity”) problem is widely recognized, and consequently there have been many attempts to improve upon standard simulation protocols. Methods which aim to generate a canonical distribution of conformations include multiple time-step methods,^{15,16} nonlinear variable transformations,¹⁷ J-walking,¹⁸ inverse renormalization group approaches,¹⁹ and adaptive resolution methods.²⁰ The most widely used class of methods, however, comprises the generalized ensemble approaches.^{21–23} Of the generalized ensemble approaches, perhaps the simplest and most popular is parallel tempering, in which a number of copies of the system are evolved in parallel at different temperatures.^{24–27} Occasionally, configurations are swapped between neighboring replicas, presumably allowing the low temperature replica to access more configuration space via high temperature conformations.

Numerous,^{28–33} as well as extensive^{34,35} parallel tempering simulations have been published, including some which claim to demonstrate the superior efficiency of the method over standard molecular dynamics (MD) simulation.^{36–38} Regardless of the validity of those claims, there appears to be a limit to the utility of parallel tempering for the study of large proteins, nucleic acids, and macromolecular complexes: the number of replicas required to bridge a specified temperature gap increases as the square root of the number of degrees of freedom of the solution.^{39,40} In other words, if atomic resolution information is desired, then very many atomic resolution simulations are required. Recent work by Berne and coworkers partly addresses this problem for explicitly solvated systems, so that the number of replicas scales with the number of degrees of freedom of the solute only.⁴¹ Solutes like proteins, of course, can be quite large.

In this paper, we address the problem of insufficient sampling of implicitly solvated biomolecules using a different approach. We recently developed an algorithm, called resolution exchange, which uses a distribution generated by a coarse-grained model to improve the sampling of a higher-resolution simulation.¹ The resolution exchange (ResEx) algorithm guarantees canon-

ical sampling for each level of resolution, so that the coarse-grained simulation introduces no bias into the high-resolution simulation. The algorithm is similar in spirit to other exchange simulations, in that conformations are swapped between otherwise independent simulations. However, by employing replicas of reduced resolution, ResEx has the potential for significant efficiency gains. Other authors have recognized the value of improving sampling with reduced resolution representations. For example, Iftimie *et. al.* used a classical potential as an importance function to improve sampling of an *ab initio* potential.⁴² Here, our goal is to sample a *classical* atomic resolution potential, which leads us to a different algorithm. Also, Liu and Sabatti formally introduced a Markov chain Monte Carlo method which allows jumps between spaces of different dimensions.⁴³ Their algorithm has apparently not been applied to the simulation of macromolecules. Lwin and Luo recently introduced an algorithm similar to ours, but it does not generate canonical sampling.⁴⁴

We have also employed a modification of the usual parallel protocol used to carry out exchange simulations,¹ generalizing the “J-walking” approach previously introduced by Frantz *et. al.*¹⁸ The J-walking (or as we call it, “top-down” exchange) method allows an unequal distribution of CPU time among the various replicas. We emphasize that any exchange simulation may be run in top-down mode. In contrast with other exchange methods, top-down exchange allows very little simulation time to be spent on the computationally expensive, high-resolution (or low temperature) replicas. Substantial improvement in sampling efficiency is therefore possible, in principle.

We previously applied the resolution exchange algorithm to butane and dileucine peptide.¹ Here, we confront issues which arise in the study of larger molecules. We show that a molecule may be coarsened incrementally, so that the overlap between models of neighboring resolution may be adjusted for improved sampling efficiency. We also demonstrate that resolution and temperature exchange are easily combined in a single simulation, so that sampling may be improved by both increasing temperature and decreasing resolution simultaneously. The incremental coarsening procedure is first demonstrated on dileucine, where we check that the correct conformational distribution is attained. We then demonstrate successful exchange between an all-atom model and an united-atom model of met-enkephalin, using two different exchange ladders: a ladder of varying resolution only, and a ladder which combines resolution and temperature changes. We will finish with a discussion of the next logical steps toward larger peptides and proteins.

1 Theory and methods

The results presented in this paper concern two distinct, recently introduced simulation methods.¹ The first is resolution exchange, which allows exchange between simulations at different resolutions, and preserves canonical sampling. The second is top-down exchange, which allows unequal distribution of CPU time, maximizing the efficiency of an exchange simulation. In addition, we describe a general incremental coarsening strategy for building a ladder of models which improves exchange efficiency.

1.1 Resolution exchange

Resolution exchange (ResEx) is motivated by the effectiveness of coarse-grained models for sampling of protein conformations,^{45,46} and by the need for atomic-level resolution for many calculations. Res-ex uses coarse-grained simulation to accelerate basin-hopping in more detailed models. In contrast with *ad hoc* methods, ResEx guarantees canonical sampling of the atomic-resolution model.

The basic idea, as in any exchange simulation, is to exchange conformations between two simulations. How are trial configurations constructed for an exchange between models with different numbers of degrees of freedom? Consider a pair of models: a coarse-grained model, a configuration of which is described by a set of coordinates Φ , and an atomic resolution model described by a larger set, $\{\Phi, \mathbf{x}\}$. Note that the coarse model is built from a *subset* of the coordinates of the detailed model. Before the exchange, let the coarse-grained configuration be Φ_a , and let the atomic-resolution coordinates be $\{\Phi_b, \mathbf{x}_b\}$. By swapping only coarse variables, the trial configuration for the coarse-grained model is simply Φ_b , and for the atomic-resolution model is $\{\Phi_a, \mathbf{x}_b\}$.

The exchange criterion is derived by considering the two simulations to constitute a single system characterized by the combined coordinates $\{\Phi_a, (\Phi_b, \mathbf{x}_b)\}$. Because the simulations—aside from the exchanges—run independently, the probability distribution of the combined system is the simple product of the individual distributions. Let the potential functions of the high- and low-resolution simulations be $U_H(\Phi, \mathbf{x})$ and $U_L(\Phi)$ respectively, and denote the Boltzmann factors as $\pi_H(\Phi, \mathbf{x}; \beta_H) = e^{-\beta_H U_H(\Phi, \mathbf{x})}/Z_H$ and $\pi_L(\Phi; \beta_L) = e^{-\beta_L U_L(\Phi)}/Z_L$, where Z_H and Z_L are the partition functions. Then the exchange attempt is accepted with the Metropolis rate:

$$\min \left[1, \frac{\pi_H(\Phi_a, \mathbf{x}_b; \beta_H) \pi_L(\Phi_b; \beta_L)}{\pi_H(\Phi_b, \mathbf{x}_b; \beta_H) \pi_L(\Phi_a; \beta_L)} \right]. \quad (1)$$

The dependence on inverse temperature (β) is made explicit, as a reminder that the method is naturally combined with temperature exchange, though this of course extends to any type exchange, such as Hamiltonian exchange.⁴⁷ Note that in the case of ordinary (temperature based) replica exchange, in which all the coordinates are swapped, Eq. (1) reduces to the familiar expression $\min[1, \exp(-\Delta\beta\Delta U)]$.

In a parallel implementation, Eq. (1), together with the protocol for trial move construction, ensures that the algorithm satisfies the detailed balance condition. To see this, consider “old” (o) and trial/“new” (n) configurations of the combined system. In the construction of any Boltzmann-preserving Monte Carlo move, two transition probabilities must be considered: the conditional probability $\alpha(o \rightarrow n)$ of *generating* the move from configuration o to n , and the conditional probability $w(o \rightarrow n)$ of *accepting* the move.⁴⁸ Detailed balance insists that $p(o)\alpha(o \rightarrow n)w(o \rightarrow n) = p(n)\alpha(n \rightarrow o)w(n \rightarrow o)$, where $p(j)$ is the equilibrium probability of configuration j . Yet the acceptance criterion (1) has the form

$$\frac{w(o \rightarrow n)}{w(n \rightarrow o)} = \frac{p(n)}{p(o)}, \quad (2)$$

implying that the generating probabilities α are identical. This is indeed the case: given a pre-defined division into coarse and detailed coordinates, the conditional probability for the move $o = \{\Phi_a, (\Phi_b, \mathbf{x}_b)\} \rightarrow n = \{\Phi_b, (\Phi_a, \mathbf{x}_b)\}$, and its inverse, are both one. That is, given the old configuration of the combined system, there is a unique trial configuration.

Lwin and Luo have constructed a similar algorithm, except that *before* checking acceptance via Eq. (1), the high-resolution trial configuration is *minimized*.⁴⁴ Such minimization (even subject to constraints on the coarse coordinates Φ , as in ref.⁴⁴) violates the detailed balance condition by biasing the generating probability, $\alpha(o \rightarrow n)$, without any compensating correction in the acceptance criterion. Put more simply, reverse moves into un-minimized configurations are impossible. The consequences of the violation are readily seen, as shown in Sec. 2.1.

1.2 Incremental Coarsening

An important practical issue is raised, however, by the construction of trial moves *without* minimization. The problem is that the degrees of freedom in the high resolution simulation $\{\Phi, \mathbf{x}\}$ are strongly coupled—for a protein, think of Φ as backbone degrees of freedom (DoF) and \mathbf{x} as side-chain DoF. Then it is clear that construction of trial moves by our method may lead

to high rejection rates. We have solved this problem by noting that the rejection rate depends on the both the *number and type* of DoF in the set $\{\mathbf{x}\}$. Employing a ladder of incremental models at intermediate resolutions allows the acceptance rates to be tuned to reasonable values, as shown in Fig. 1.

A ladder of incrementally coarsened models is straightforward to construct. Consider coarsening from an all-atom representation of a protein to a united-atom representation. In the first model above the all-atom level, only one residue is described by the united-atom representation—the protein is described by a “mixed model”, with one united-atom residue and the rest all-atom. Then, in the next level up, there are two united-atom residues, and so on, until the entire protein is described by the united-atom force field. A similar procedure may then be used to go beyond the united-atom level to a united *residue* level. Notice that it may be desirable to coarsen more than one residue at a time, since some residues have fewer degrees of freedom than others.

Of course, implementation of the incremental ladder just described requires the construction of a potential function which has both united- and all-atom groups. In this work, we have built this mixed potential by combining the parameters for united and all-atom force fields into a single file. In other words, we created a larger parameter file, which contains both all-atom and united atom atom types. This file also includes all of the interactions for both united- and all-atom types, with the united-atom interactions modified as described in Sec. 1.1. Adding some parameters (taken from the all-atom potential) for the interfaces, where united and all-atom residues link, the mixed potential describes the whole molecule. The parameters (formatted for use in TINKER) are included as supplementary material.

The incremental coarsening approach just described is rather general and not restricted to implementing exchange ladders spanning united- to all-atom resolutions. Lower resolution models could also be considered, for which it may be desirable to coarsen several residues at once. A first quantitative analysis of the incremental coarsening procedure, suggesting how efficiency can be improved, is given in Sec. 2.2 and 2.3.

1.3 Top-down exchange

In many exchange simulations, whether they are temperature-based,³⁸ Hamiltonian-based,⁴⁷ or use some other extended ensemble,^{49,50} the goal is to improve the sampling of a hard-to-sample model (such as an all-atom protein model at native conditions) by sampling a related model, which is *presumed* easier-

to-sample (such as the same all-atom model at higher temperature)¹. Information is swapped between the simulations by occasionally exchanging configurations, in a way which preserves canonical sampling of each distribution. Usually there is little overlap between the hard-to-sample (henceforth, “bottom level”) and the easy-to-sample (henceforth, “top-level”) models, and therefore a ladder of intermediate models is required.

A critical observation is that the accuracy which is ultimately attained in the hard-to-sample, bottom-level model is effectively limited by that which is obtained in the easy-to-sample, top-level model.¹⁸ In many cases, the top level is still quite difficult to sample well, and will require considerable CPU time to reach an acceptable accuracy—much more than it would usually be allotted in a parallel implementation. It is this observation which motivates the top-down method. The top-down algorithm shown schematically in Fig. 2 was developed previously for temperature-based simulation,¹⁸ though was not widely recognized as such. The procedure is as follows:

- (i) Run and store a simulation at the top level (model M_N) until it is judged to be sufficiently converged. This trajectory is a sample of the distribution $\pi_N(\mathbf{r})$ of the top level, where N labels the level and \mathbf{r} labels the configurations. In the case of ResEx, $\mathbf{r} = (\Phi, \mathbf{x})$. Let n be a running index, with $n = N$ at this top level.
- (ii) Start a simulation at the $n - 1$ level—for example, at the next lower temperature. Configurations \mathbf{r}_{n-1} will be sampled according to π_{n-1} for model M_{n-1} .
- (iii) Whenever an exchange is to be attempted, pull a random trial configuration \mathbf{r}_n from the M_n trajectory. In the case of ResEx, one requires only the subset Φ .
- (iv) Accept the trial configuration according to

$$\min \left[1, \frac{\pi_{n-1}(\mathbf{r}_n) \pi_n(\mathbf{r}_{n-1})}{\pi_{n-1}(\mathbf{r}_{n-1}) \pi_n(\mathbf{r}_n)} \right],$$

where $\pi_i(\mathbf{r}) = e^{-\beta_i U_i(\mathbf{r})} / Z_i$, Z_i is the partition function, U_i is the potential function, and β_i is the inverse temperature. Notice the partition functions need not be known, as they cancel between numerator and denominator.

(v) Continue with steps (iii) and (iv) until the sampling of the $n - 1$ level is judged sufficient. Store the $n - 1$ trajectory.

(vi) Continue with steps (ii) to (v) for $n = N - 2, N - 3, \dots$ until the bottom level simulation is complete.

¹For a discussion of these issues from a statistical perspective, see Neal⁵¹

First, note that canonical sampling is maintained by Eq. (1.3),¹⁸ just as in an ordinary parallel exchange simulation. On the other hand, detailed balance is not satisfied, as the trial configuration for the level n simulation (\mathbf{r}_{n-1} above) is discarded—making reverse moves effectively impossible. The error is one of practice, not of principle, arising from the fact that the samples of π_n and π_{n-1} are finite, just as in any simulation.

To see intuitively that canonical sampling is maintained by top-down exchange, imagine a pair of simulations undergoing ordinary parallel exchange. Unbeknownst to the investigator, however, the top level simulation is running on a very fast processor, while the other is running on an old, slow processor. Between neighboring exchange points, the trial conformations from the fast processor will be far more decorrelated than those of the slow simulation—which mimics the effect of the top-down protocol. However, these exchanges still satisfy detailed balance. In the limit of an infinitely fast top-level simulation, trial configurations are completely decorrelated, and one could equally well choose randomly from π_n as in step (iii).

Second, notice that because trial configurations are pulled *at random* from the sample of π_n in step (iii), transitions which are slow in the actual M_n trajectory occur rapidly among the trial configurations. Maximum benefit is thus obtained from successful exchanges—in contrast with a parallel exchange simulation, where trial configurations are typically separated by only a few picoseconds, and remain highly correlated.

Third, good results may be obtained expending very little CPU time on all levels except the top one. This may be understood from an energy landscape perspective. The top level has been used to thoroughly sample the space—high barriers are crossed, and major sub-basins equilibrated. At lower levels, only local equilibration need occur. For example, let τ_{nonloc} be the time to cross high barriers, τ_{local} be the time to equilibrate locally, and say that m successful exchanges are needed to sample the space well. Then $\tau_{local} \times m$ CPU time is needed to sample the lower level. The required condition to save time over a parallel simulation is that $\tau_{local} \ll \tau_{nonloc}$. The degree to which this condition is satisfied will depend on the system studied, but the top-down approach allows the flexibility to take advantage of a separation in time scales. This idea is reminiscent of the “dragging” of fast degrees of freedom, suggested by Neal,⁵² and the multiple time step approaches developed by Berne and coworkers.^{15,16}

Finally, a major advantage of top-down simulation over parallel exchange protocols is that exchange attempts are nearly “free”, in the sense that no communication is required between processors.¹⁸ This means that exchanges may be attempted very frequently, and therefore much lower exchange rates

may be accommodated. In the case of temperature exchange, this allows either for the steps in the temperature ladder to be more widely spaced, or for the treatment of larger systems with fewer replicas.

1.4 Simulation details

In ideal circumstances, low-resolution models used in ResEx simulations would be specifically optimized for resolution exchange. They would have maximal conformational overlap for the common degrees of freedom. Here we work with an “off the shelf” low resolution model (united atom) which leads to some difficulties. Consider, for example, a C^α - C' bond which is parameterized in the two models by two slightly different natural bond lengths. In an exchange attempt, the configurations are swapped, and in each trial configuration the C^α - C' bond is moved a bit from its preferred length. These small contributions add up for every harmonic term in the entire molecule, and have a noticeable effect on the acceptance of exchange moves. We have solved this problem by simply changing the harmonic parameters of the coarse model to match those of the detailed model. This makes the coarse model more “exchangeable” with the detailed model. Since the coarse model is simply “suggesting” configurations for the atomic model, and since Eq. (1) guarantees that no bias is introduced by the coarse model, we need not worry that the coarse model parameters are changed from their original values. We now describe in detail the two molecular systems which were studied in the present work.

Dileucine. We first studied dileucine peptide (ACE-(Leu)₂-NME) using the same forcefield parameters as in Ref.¹ Here, we also carried out an incrementally coarsened ResEx simulation of dileucine in 5 levels. The coarsest level (M_4) was the same as in,¹ namely a modified version of OPLSUA.⁵³ In lower levels, the molecule was rendered in finer detail beginning at the N-terminus: in M_3 , the N-terminal methyl group, C^α , and C^β of the first residue were modelled in full atomistic detail; in M_2 , C^γ and both C^δ 's of the first residue were additionally modelled explicitly; in M_1 , the C^α , C^β , C^γ , and one C^δ of the second residue were modelled explicitly; and finally in M_0 , the entire molecule was rendered in full atomic detail. The ladder of mixed models was chosen to keep approximately fixed the number of hydrogens by which neighboring levels differ, without splitting a methyl group.

The top level (M_4) was simulated first, for 25, 50, 100, or 200 nsec. The different lengths of top-level simulation were used to generate the different data points in Fig. 4. Then the higher resolution models were run, as per the top-down protocol (see Sec. 1.3), attempting exchanges once per psec.

A total of 2.5×10^3 exchanges attempted between each level, for a total trajectory length per level of 2.5 nsec. Frames were stored every 0.1 psec, for a total of 2.5×10^4 frames in the sample at each level below the top.

Met-enkephalin. We next studied met-enkephalin (NH_3^+ -Phe-Gly-Gly-Tyr-Met- COO^-). The united atom force field was a modified version of OPLSUA.⁵³ The force field was modified so that the bond length and angle bending parameters matched those of the all-atom force field, which improves exchangeability (or conformational overlap) of the two models. The sample of the top level model was constructed from two independent 100 nsec trajectories, both started from pdb structure 1plw (1st NMR model), generated by Langevin dynamics as implemented in TINKER v. 4.2.⁵⁴ The friction coefficient was 5 psec^{-1} , and solvation was modelled with the GB/SA method.⁵⁵ The first 1 nsec of each trajectory was discarded and frames were stored every 10 psec for a total of 19,800 frames in the sample.

We then ran the next higher resolution simulation, as per the top-down algorithm (see Sec. 1.3). This model was of mixed resolution, with the Tyr₁ residue represented by the OPLSAA all-atom forcefield,⁵⁶ and the remaining residues described by the OPLSUA force field. Every 1 psec, a random configuration was pulled from the top-level (M_5) trajectory, and a resolution exchange was attempted, with acceptance governed by Eq. (1). Since the acceptance ratio for the M_5 to M_4 exchanges was approximately 10%, the average length of M_4 trajectory between exchanges was 10 psec. A total of 10^4 ResEx moves were attempted, for a total M_4 trajectory length of 10 nsec. Frames were stored every 0.1 psec for a sample of 10^5 frames.

This procedure was then repeated for each level shown in Fig. 1, with the exception that the attempt frequency of ResEx moves was adjusted for the acceptance ratios, so that the segments of the simulations between exchanges were kept approximately constant at 10 psec. Also, the total number of attempted exchanges was adjusted so that approximately 10^3 successful exchanges were observed between each level, for a total trajectory length of 10 nsec at each level. Given that the top level is presumed to be well-sampled, 10^3 exchanges should sample a large number of basins.

2 Results

In a previous short paper, we tested the ResEx algorithm on two small molecules: butane and dileucine peptide.¹ It was shown that the method produced results in agreement with those obtained by standard simulation methods. For the sake of clarity, here we first demonstrate our approach

on a two-dimensional toy model, consisting of two basins which differ only *entropically*. We also extend the method to two peptides, dileucine and met-enkephalin, in order to demonstrate the viability of incremental coarsening.

2.1 Results: Two-dimensional model

An important consideration in designing any sampling method is whether it will correctly account for entropy differences. We therefore designed the potential surface shown in fig. 3 to compare three different sampling methods: a “standard” Monte Carlo simulation, the same Monte Carlo with resolution exchange, and the same Monte Carlo with the “dual REM” method of Lwin and Luo.⁴⁴

The surface $U(x, y)$ in fig. 3 is described by the function

$$U(x, y) = E_b (x^2 - 1)^2 + \frac{E_0 y^2}{1 + w (\tanh(x/0.1) + 1) / 2}, \quad (3)$$

where $E_0 \equiv k_B T$, $E_b = 10 k_B T$ is the barrier height, and w controls the width of the right well in the figure. Notice that the profile of the surface at $y = 0$ is symmetric about $x = 0$: $U_x(x) \equiv U(x, y = 0) = E_b(x^2 - 1)^2$, i.e., the two minima are of equal depth. The parameters were chosen so that the equilibrium populations of the two wells differ greatly—we used $w = 500$, so that the right well holds 95% of the population, as measured by standard techniques.

For both the ResEx and the dual REM simulations, the “coarse-grained” potential was simply the one-dimensional potential U_x , i.e., a symmetric double well.

To describe the exchange moves explicitly, we denote $2D$ configurations by (x, y) and $1D$ configurations by \tilde{x} . For both algorithms, an exchange move consists of two parts: the construction of a $1D$ trial configuration (\tilde{x}_{new}) from a $2D$ configuration (x_{old}, y_{old}) , and vice versa: the construction of a $2D$ trial configuration (x_{new}, y_{new}) from a $1D$ configuration (\tilde{x}_{old}) . The construction of a $1D$ configuration in each case is simple—the “extra” (y) coordinate is dropped, i.e., $\tilde{x}_{new} = x_{old}$.

The only difference between the two simulations is in the construction of trial configurations for the $2D$ model from the $1D$ model. In ResEx, the trial configuration is the \tilde{x} coordinate from the $1D$ model, with the (old) y coordinate from the $2D$ model, i.e., $x_{new} = \tilde{x}_{old}$ and $y_{new} = y_{old}$. In dual REM, on the other hand, the trial y coordinate is chosen randomly, and then minimized. For the potential $U(x, y)$, this means that $y_{new} = 0$ *always*, i.e., $x_{new} = \tilde{x}_{old}$ and $y_{new} = 0$.

The ResEx simulation correctly samples the two wells, giving a population in the right well of $96.4 \pm 1.6\%$. The dual REM simulation, on the other hand, yields a population of $51.0 \pm 1.6\%$ for the right well. What causes the error in the dual REM simulation? The answer is that the construction of dual REM trial moves violates detailed balance. More specifically, the minimization of the y coordinate means that the difference in *width* between the two wells is not accounted for correctly, since in dual REM $y_{new} = 0$ *always*. Notice that it is not the random selection of the y coordinate which intrinsically violates detailed balance, only the subsequent minimization.

What is the analogous situation in molecular simulations? In this case, both ResEx and dual REM construct trial moves in internal coordinates—the coarse-grained model is built from a subset of the degrees of freedom of the atomic model. For example, the coarse-grained model (the x coordinate above) may be the backbone coordinates of a protein, and the remainder (the y coordinate above) may be the sidechain degrees of freedom. In dual REM construction of trial moves, the sidechains are minimized prior to exchange, and the therefore differences in entropy between different sidechain conformations are neglected. In ResEx, there is no minimization prior to exchange, and canonical sampling is maintained.

2.2 Results: Incremental coarsening of dileucine.

We previously reported on ResEx results for dileucine, demonstrating successful exchange and significant speedup from a direct exchange between all-atom and united-atom models. Here, show that dileucine may be coarsened incrementally, and that (i) the correct distribution is observed for the all-atom model and (ii) adding additional intermediate levels of resolution improves efficiency.

The additional levels boost the exchange acceptance by two orders of magnitude: exchanges were successful between M_4 and M_3 15.5% of the time, between M_3 and M_2 12.7% of the time, between M_2 and M_1 29.0% of the time, and between M_1 and M_0 44.0% of the time. By comparison, exchanging AA and UA dileucine in a single step is successful only 0.16% of the time.¹ However, we need to ask whether it is really more efficient to introduce additional levels of simulation in order to boost the acceptance of exchange moves.

In fact, it appears to be substantially more efficient to use incremental coarsening rather than abrupt coarsening. The cost for a given ladder of N

levels may be written

$$\text{total cost} = \text{top level cost} + \sum_{i=0}^{N-1} \frac{m\tau_i}{r_i}, \quad (4)$$

where the cost of the top level is fixed, m is the fixed number of successful exchanges which are desired, τ_i is the simulation cost for an interval between two exchange attempts at level i , and r_i is the acceptance rate between levels i and $i + 1$. We have assumed that the sampling of level i demands a fixed number of successful resolution exchanges, consistent with the motivation of the top-down protocol discussed in Sec. 1.3.

Eq. (4) implies that the effective exchange rate for an incremental ladder is a reciprocal sum of the individual rates. If we assume the τ_i are equal for all levels (which is exact for temperature exchange), then

$$1/r_{\text{eff}} = \sum_i 1/r_i, \quad (5)$$

giving an effective rate for the 5 level dileucine ladder of 5.1%. This result suggests an improvement in efficiency relative to the single step ladder, where the rate was 0.156%.¹

In Fig. 4, we compare the sampling of dileucine by three different simulation protocols: standard Langevin dynamics, resolution exchange with two levels, and resolution exchange with five levels. Sampling is assessed by examining the relative populations of the α and β states ($e^{-\Delta G_{\alpha\beta}/k_B T}$) considered in Ref.¹ The convergence of this relative population measure requires transitions between the two states, which occur infrequently in a standard simulation. The five-level ladder clearly outperforms the two-level ladder, as we are able to generate results both more accurate and more precise with the five-level ladder in an equal amount of CPU time. Note that the total simulation time required for the entire ladder, including the top level, is included in the ResEx data points.

The efficacy of the ResEx approach is underscored by the fact that, at the top level (united atom), the sign of $\Delta G_{\alpha\beta}$ is wrong. That is, the exchange process corrects for a substantial bias in the coarse model.

2.3 Results: Incremental coarsening of met-enkephalin

Met-enkephalin is a flexible neurotransmitter which participates in immune responses and pain inhibition, among other roles.^{57,58} By virtue of its small

size and biological interest, it often is used to test new simulation methods^{27,59,60} and compare existing protocols.^{58,61}

Using met-enkephalin, we demonstrate the efficacy of the incremental coarsening procedure for a ladder of decreasing resolution at constant temperature, and for a ladder of simultaneously decreasing resolution and increasing temperature. Because quantifying the quality of sampling for met-enkephalin is considerably more difficult than is commonly appreciated, we will not present a detailed efficiency analysis. More will be said on this second topic in the discussion.

We employed the ResEx algorithm in a top-down framework, as sketched in Fig. 2. First, the top-level simulation (coarsest resolution—here, united-atom) was run. We then ran an exchange simulation at the next highest resolution—here, one residue was represented at all-atom resolution, and the rest of the peptide was united-atom. This procedure was continued, “de-coarsening” one residue at a time, until the entire peptide was represented at the all-atom level. Details are given in Secs. 1.2 and 1.4.

The incremental coarsening procedure substantially increases exchange acceptance. The five rates in the six-level ladder vary from 2.4% to 18%, as shown in Fig. 1. For comparison, exchanging between all-atom and united-atom models of met-enkephalin, with no intermediate levels of resolution, results in an acceptance ratio of 0.09%. The acceptance ratios vary, in part, according to the number of degrees of freedom by which the two levels differ.

For met-enkephalin, the principal results are the acceptance rates shown in Fig. 1, which are significant for several reasons. First, they demonstrate the first implementation of the incremental coarsening approach in a complex peptide. Second, because they are well within the practical range of the top-down protocol—see Sec. 1.3 and the Discussion—they indicate that the ResEx algorithm could prove important for larger peptides. Lastly, by comparing the effective exchange rate suggested by Eq. (5), $r_{\text{eff}} = 5.2\%$, with the rate of 0.09% for direct exchanges between united- and all-atom models, one sees that a substantial speedup has been achieved. Of course, the magnitude of the improvement is merely suggestive—without a rigorous quantification of the sampling quality, there can be no rigorous comparison of efficiency. Such a quantification is beyond the scope of this work, as noted in the Discussion.

It is useful to understand the intuitive reason behind the advantage of incremental coarsening. If one writes the acceptance criteria (1) and (1.3) in the form $\min[1, e^{-\epsilon}]$, then for exchanges between models of greatly differing resolution, one typically finds the dimensionless “energy” is large, i.e. $\epsilon \gg 1$. It seems to be roughly true that this energy is proportional to the differ-

ence in the number of degrees of freedom in the models being exchanged. However, if the change is made incrementally using many models M_i , then between levels i and $i+1$ there is a relatively small cost $\Delta\epsilon_i$, with $\sum_i \Delta\epsilon_i \sim \epsilon$. It is clear that with enough increments, one can achieve $\Delta\epsilon_i \ll 1$, and thus create a high likelihood for exchange since the corresponding Boltzmann factors are much larger: $r_i \sim e^{-\Delta\epsilon_i} \gg e^{-\epsilon}$. This is exactly what is embodied in Eq. (4). The trade-off is that one pays the cost for simulating the additional intermediate levels. However, as has been stressed in Sec. 1.3, the intermediate-level simulations are very short compared to the top level. In the present context, for instance, the top level met-enkephalin trajectory is 198 nsec, while all other levels were simulated for only 10 nsec. The net savings can be quite substantial, especially considering that good sampling is achieved by increasing the number of exchanges.

While we cannot yet rigorously measure sampling quality, we can show that the results obtained with ResEx are consistent with those obtained by standard methods, by comparing Ramachandran histograms (Fig. 5) from the ResEx simulation, to those obtained by standard simulation (990 nsec of simulation with the M_0 parameters). Overall, the agreement between the ResEx simulation and the 990 nsec conventional simulation is quite good. However, a careful comparison reveals a region on the Phe₄ plot, labelled “A”, which is noticeably under-sampled by the ResEx simulation, as compared to the 990 nsec Langevin dynamics trajectory. The explanation is provided by an inspection of the Phe₄ histogram of the united-atom simulation: region “A” was not sampled by the top-level simulation. The failure points to a potential weakness of the ResEx (or any exchange) method—regions which are not sampled by the top level will be difficult to sample in any of the other levels. This is a specific instance of a general problem that occurs whenever auxiliary distributions are used to enhance sampling of some “distribution of interest,” namely the need to balance overlap with wider sampling via the auxiliary distribution.⁵¹ In other words, it is a failure of the top-level simulation rather than the algorithm.

Interestingly, Fig. 5 also presents two counterexamples to the forgoing discussion. Regions “B” of the Gly₃ and “C” of the Met₅ plots were both well-sampled by the ResEx simulation, despite being infrequently visited by the top-level. That is, the ResEx acceptance criterion (1) correctly “re-weights” the conformation space of the all-atom model by allowing normal dynamics to continue when appropriate. Nevertheless, we are in the process of experimenting with other “schedules” (combinations of attempt frequency and number of exchange attempts) to balance the normal and the exchange dynamics.

Ideally, the coarse model distribution would have better overlap with the high-resolution distribution, and the balance could be adjusted to favor exchanges over normal dynamics. This would allow the same quality of sampling with less simulation at each level below the top. In the long term, we hope to design coarse models constructed to *not* eliminate *any* regions of configuration space in more detailed models.

2.4 Resolution exchange with tempering

We have also explored the possibility of combining resolution exchange with parallel tempering, so that the sampling of the reduced models is improved both by the reduction in resolution and by increased temperatures. In a standard parallel tempering simulation, the temperatures are roughly exponentially distributed, in order that the conformational overlap between neighboring temperatures is constant over the ladder. However, there is no simple relationship between the change in *resolution* and the acceptance of resolution exchanges. Some care must therefore be taken with the assignment of the temperature ladder.

We began with the ladder of models in fig. 1. The acceptance ratios give an idea of the temperature gap which may be tolerated between two levels—a higher acceptance ratio will tolerate a larger jump in temperature. However, compared to standard parallel tempering simulations,^{26–35} it may seem that the acceptance ratios are already too low to accommodate tempering in addition to resolution exchange. After all, we may expect that any difference in temperature will lower the acceptance of exchange moves. In this regard, the top-down approach has an important advantage over a parallel implementation. Since exchange attempts are “free” (no communication between processors is required), they may be attempted much more frequently, and lower acceptance ratios may be tolerated.¹⁸ Indeed, in our original study of dileucine peptide with top-down resolution exchange, the acceptance ratio was much less than 1%.¹ See also Sec. 1.3.

The ladder combining temperature and resolution is shown in Fig. 6. The temperature gaps were chosen by trial and error, aiming for an acceptance of attempted exchanges of a few percent between neighboring levels. Based upon this restriction, the top-level simulation was run at a temperature of 700 K, which is comparable to previously published parallel tempering studies of met-enkephalin.^{27,36} We should expect, however, that fixed CPU cost sampling should be improved relative to ordinary replica exchange, by virtue of the reduction in resolution.

The reduction in resolution confers an additional benefit when combined

with tempering. Since the overlap between neighboring levels in a parallel tempering simulation scales like (number of DoF)^{1/2}, reducing resolution allows the temperature gaps to *increase* as the resolution is reduced. The overlap between neighboring levels in a combined resolution/tempering ladder is thus controlled both by the change in resolution, and the change in temperature, with the two effects compensating one another in an unknown way. Indeed, we observed one puzzling case in our search for an appropriate resolution/temperature ladder. In one ladder (data not shown), exchange between levels M₂ and M₃ was successful about 7% of the time when both were at 298 K, while exchange occurred approximately 11% of the time when M₂ was thermostatted to 305 K, and M₃ to 320 K. We have not explained this result—though it should be remembered that different models have different landscapes, and therefore temperatures may not be directly compared.

3 Concluding Discussion

We have extended our resolution exchange (ResEx) method¹ using an incremental coarsening procedure for implicitly solvated peptides. After carefully testing the approach in the two-residue dileucine peptide, we applied it successfully to the five-residue met-enkephalin. Incremental coarsening allows tuning of the conformational overlap between models of differing resolution, and therefore makes practical simulations which would otherwise be hampered by poor acceptance of exchange moves. We have also demonstrated that resolution exchange is naturally combined with parallel tempering, so that the reduced resolution models may be aided in their sampling of conformations by elevated temperatures.

Ramachandran histograms demonstrate that, for the most part, ResEx simulation is consistent with standard simulation techniques. In one case, however, they reveal a weakness of our method—a top-level simulation which eliminates important regions of conformation space will result in poor sampling at the bottom level. This weakness is shared by any simulation which relies upon auxiliary ensembles to sample among major sub-basins. In the future we hope to eliminate this problem by more careful construction of reduced models.

Of course, we hope to treat still larger molecules with the ResEx method. Since it is essential that the top-level be well-sampled, the treatment of larger molecules will require yet coarser top-level simulation. This will likely require incremental coarsening from the united-atom level to a model with

one or two beads per residue. Suitable models are under development. It appears that the ResEx approach cannot be applied easily to explicitly solvated systems; however, given the difficulty and importance of sampling implicitly solvated systems, ResEx may prove very valuable for biomolecular simulation.

We have also developed an alternative rigorous algorithm which permits the use of coarse top-level simulations to generate atomically detailed canonical samples. It is essentially a “decorating” procedure, and it eliminates the potential issue of correlations between coarse coordinates Φ and detailed coordinates \mathbf{x} , which could reduce acceptance rates in resolution exchange. Specifically, after generating a low-resolution sample distributed according to $\pi_L(\Phi)$, one can independently sample detailed coordinates \mathbf{x} according to an arbitrary distribution $\pi_x(\mathbf{x})$. (For example, π_x could be based on harmonic terms in the full forcefield.) Full configurations are thus generated according to the simple product $\pi_L(\Phi)\pi_x(\mathbf{x})$ and may be re-weighted to generate a fully detailed, high-resolution, distribution $\pi_H(\Phi, \mathbf{x})$ using standard methods.⁶² In the long term, the decorating approach may prove useful for adding explicit solvent. It may also be implemented in an incremental fashion.

An “auxiliary” question which remains to be carefully addressed is the quantification of sampling efficiency. There are numerous proposals for judging whether a simulation is converged—some are based on principal components,⁶³ others on energy-based ergodic measures,⁶⁴ and our own work in progress uses structural histograms.⁶⁵ Which one provides an appropriate measure depends on what properties are of interest. For applications which depend on the relative populations of various conformations, such as calculation of binding affinities for small molecules, a measure which depends directly on the conformational distribution is needed. Such a method is under development—for now we only mention that structural histograms provide a much more sensitive signal of non-convergence than energy-based methods.⁶⁵

Acknowledgements. The authors would like to thank Marty Ytreberg, Bruce Berne, Michael Deem, and Angel García for interesting discussions and insightful comments. This work was supported by the Department of Environmental and Occupational Health and the Department of Computational Biology at the University of Pittsburgh, and by the NIH (grants ES007318 and GM070987).

Supporting Information. Parameter files for the modified oplsua force field and for the mixed ua-aa force field are provided. The files are formatted for use with the Tinker molecular modelling simulation package. Also

included is a sample '.key' file for running the mixed forcefield simulations in Tinker. This information is available free of charge via the internet at <http://pubs.acs.org>.

References

- (1) Lyman, E.; Ytreberg, F. M.; Zuckerman, D. M. *Phys. Rev. Lett.* **2006**, *96*, 028105.
- (2) Aksimentiev, A.; Balabin, I. A.; Fillingame, R. H.; Schulten, K. *Biophys. J.* **2004**, *86*, 1332–1344.
- (3) Duan, Y.; Kollman, P. A. *Science* **1998**, *282*, 740–744.
- (4) Snow, C. D.; Nguyen, H.; Pande, V. S.; Gruebele, M. *Nature* **2002**, *420*, 102–106.
- (5) Halgren, T. A.; Damm, W. *Curr. Op. Struct. Bio.* **2001**, *11*, 236–242.
- (6) Anisimov, V. M.; Lamoureux, G.; Vorobyov, Igor, V.; Huang, N.; Roux, B.; MacKerell Jr., A. D. *J. Chem. Theor. Comp.* **2005**, *1*, 153–168.
- (7) Iftimie, R.; Minary, P.; Tuckerman, M. E. *Proc. Nat. Acad. Sci. USA* **2005**, *102*, 6654–6659.
- (8) Huang, Y. J.; Montelione, G. T. *Nature* **2005**, *438*, 36–37.
- (9) McCallum, S. A.; Hitchens, T. K.; Torborg, C.; Rule, G. S. *Biochemistry* **2000**, *39*, 7343–7356.
- (10) Volkman, B. F.; Lipson, D.; Wemmer, D. E.; Kern, D. *Science* **2001**, *291*, 2429–2433.
- (11) Shoichet, B. K. *Nature* **2004**, *432*, 862–865.
- (12) Comeau, S.; Vajda, S.; Camacho, C. J. *PROTEINS* **2005**, *60*, 239–244.
- (13) Kollman, P. *Chem. Rev.* **1993**, *93*, 2395–2417.
- (14) Roderinger, T.; Pomès, R. *Curr. Op. Struct. Bio.* **2005**, *15*, 164–170.
- (15) Humphreys, D. D.; Friesner, R. A.; Berne, B. J. *J. Phys. Chem.* **1994**, *98*, 6885–6892.
- (16) Hetényi, B.; Barnacki, K.; Berne, B. J. *J. Chem. Phys.* **2002**, *117*, 8203–8207.

- (17) Zhu, Z.; Tuckerman, M. E.; Samuelson, S. O.; Martyna, G. J. *Phys. Rev. Lett.* **2002**, *88*, 100201-1–100201-4.
- (18) Frantz, D. D.; Freeman, D. L.; Doll, J. D. *J. Chem. Phys.* **1990**, *93*, 2768–2783.
- (19) Bai, D.; Brandt, A. “NATO science series III”, Technical Report 177, 2001.
- (20) Praprotnik, M.; Site, L. D.; Kremer, K. *J. Chem. Phys.* **2005**, *123*, 224106.
- (21) Mitsutake, A.; Sugita, Y.; Okamoto, Y. *Biopolymers* **2001**, *60*, 96–123.
- (22) Gnanakaran, S.; Nymeyer, H.; Portman, J.; Sanbonmatsu, K. Y.; García, A. E. *Curr. Op. Struct. Bio.* **2003**, *13*, 168–174.
- (23) Earl, D. J.; Deem, M. W. *Phys. Chem. Chem. Phys.* **2005**, Advance article.
- (24) Swendsen, R. H.; Wang, J.-S. *Phys. Rev. Lett.* **1986**, *57*, 2607–2609.
- (25) Hukushima, K.; Nemoto, K. *J. Phys. Soc. Jpn.* **1996**, *65*, 1604.
- (26) Hansmann, U. H. E. *Chem. Phys. Lett.* **1997**, *281*, 140–150.
- (27) Sugita, Y.; Okamoto, Y. *Chem. Phys. Lett.* **1999**, *314*, 141–151.
- (28) García, A. E.; Sanbonmatsu, K. Y. *PROTEINS* **2001**, *42*, 345–354.
- (29) Zhou, R.; Berne, B. J. *Proc. Nat. Acad. Sci. USA* **2002**, *99*, 12777–12782.
- (30) Im, W.; Feig, M.; Brooks, C. L. I. *Biophys. J.* **2003**, *85*, 2900–2918.
- (31) Kokubo, H.; Okamoto, Y. *J. Chem. Phys.* **2004**, *120*, 10837–10847.
- (32) Im, W.; Brooks, C. L. I. *J. Mol. Biol.* **2004**, *337*, 513–519.
- (33) Gnanakaran, S.; Hochstrasser, R. M.; García, A. E. *Proc. Nat. Acad. Sci. USA* **2004**, *101*, 9229–9234.
- (34) García, A. E.; Onuchic, J. N. *Proc. Nat. Acad. Sci. USA* **2003**, *100*, 13898–13903.

- (35) Paschek, D.; Gnanakaran, S.; García, A. E. *Proc. Nat. Acad. Sci. USA* **2005**, *102*, 6765–6770.
- (36) Sanbonmatsu, K. Y.; García, A. E. *PROTEINS* **2002**, *46*, 225–234.
- (37) Roe, D. R.; Hornak, V.; Simmerling, C. *J. Mol. Biol.* **2005**, *352*, 370–381.
- (38) Zhang, W.; Wu, C.; Duan, Y. *J. Chem. Phys.* **2005**, *123*, 154105-1–154105-9.
- (39) Kofke, D. A. *J. Chem. Phys.* **2002**, *117*, 6911–6914.
- (40) Kone, A.; Kofke, D. A. *J. Chem. Phys.* **2005**, *122*, 206101-1–206101-2.
- (41) Liu, P.; Kim, B.; Friesner, R. A.; Berne, B. J. *Proc. Nat. Acad. Sci. USA* **2005**, *102*, 13749–13754.
- (42) Iftimie, R.; Salahub, D.; Wei, D.; Schofield, J. *J. Chem. Phys.* **2000**, *113*, 4852–4862.
- (43) Liu, J. S.; Sabatti, C. Simulated sintering: Markov chain Monte Carlo with spaces of varying dimensions. In *Bayesian Statistics*, Vol. 6; Bernardo, J.; Berger, J.; Dawid, A.; Smith, A., Eds.; Oxford university press, New York: 1998.
- (44) Lwin, T. Z.; Luo, R. *J. Chem. Phys.* **2005**, *123*, 194904-1–194901-9.
- (45) Levitt, M.; Warshel, A. *Nature* **1975**, *253*, 694–698.
- (46) Bradley, P.; Misura, K. M. S.; Baker, D. *Science* **2005**, *309*, 1868–1871.
- (47) Fukunishi, H.; Watanabe, O.; Takada, S. *J. Chem. Phys.* **2002**, *116*, 9058–9067.
- (48) Frenkel, D.; Smit, B. *Understanding Molecular Simulation*; Academic Press: San Diego, 1996.
- (49) Yan, Q.; de Pablo, J. J. *J. Chem. Phys.* **2000**, *113*, 1276–1282.
- (50) Sugita, Y.; Kitao, A.; Okamoto, Y. *J. Chem. Phys.* **2000**, *113*, 6042–6050.
- (51) Neal, R. M. **2005**, <http://xxx.lanl.gov/abs/math.ST/0511216>.

- (52) Neal, R. M. “Taking bigger Metropolis steps by dragging fast variables”, Technical report No. 0411, Dept of Statistics, University of Toronto, 2004 <http://xxx.lanl.gov/abs/math.ST/0502099>.
- (53) Jorgensen, W. L.; Madura, J. D.; Swenson, C. J. *J. Am. Chem. Soc.* **1984**, *106*, 6638–6646.
- (54) Ponder, J. W.; Richard, F. M. *J. Comput. Chem.* **1987**, *8*, 1016–1024 <http://dasher.wustl.edu/tinker/>.
- (55) Still, W. C.; Tempczyk, A.; Hawley, R. C. *J. Am. Chem. Soc.* **1990**, *112*, 6127–6129.
- (56) Jorgensen, W. L.; Maxwell, D. S.; Tirado-Rives, J. *J. Am. Chem. Soc.* **1996**, *117*, 11225–11236.
- (57) Plotnikoff, N. P.; Faith, R. E.; Murgo, A. J.; Good, R. A. *Enkephalins and endorphins: stress and the immune system*; Plenum: New York, 1986.
- (58) Shen, M.-Y.; Freed, K. F. *Biophys. J.* **2002**, *82*, 1791–1808.
- (59) Mitsutake, A.; Sugita, Y.; Okamoto, Y. *J. Chem. Phys.* **2003**, *118*, 6664–6675.
- (60) Mitsutake, A.; Kinoshita, M.; Okamoto, Y.; Hirata, F. *J. Phys. Chem. B* **2004**, *108*, 19002–19012.
- (61) Zaman, M. H.; Shen, M.-Y.; Berry, R. S.; Freed, K. F. *J. Phys. Chem.* **2003**, *107*, 1686–1691.
- (62) Ferrenberg, A. M.; Swendsen, R. H. *Phys. Rev. Lett.* **1988**, *61*, 2635–2638.
- (63) Hess, B. *Phys. Rev.* **2002**, *E65*, 031910-1–031910-10.
- (64) Straub, J. E.; Rashkin, A. B.; Thirumalai, D. *J. Am. Chem. Soc.* **1994**, *116*, 2049–2063.
- (65) Lyman, E.; Zuckerman, D. M. Preprint: <http://xxx.lanl.gov/abs/physics/0601104>.

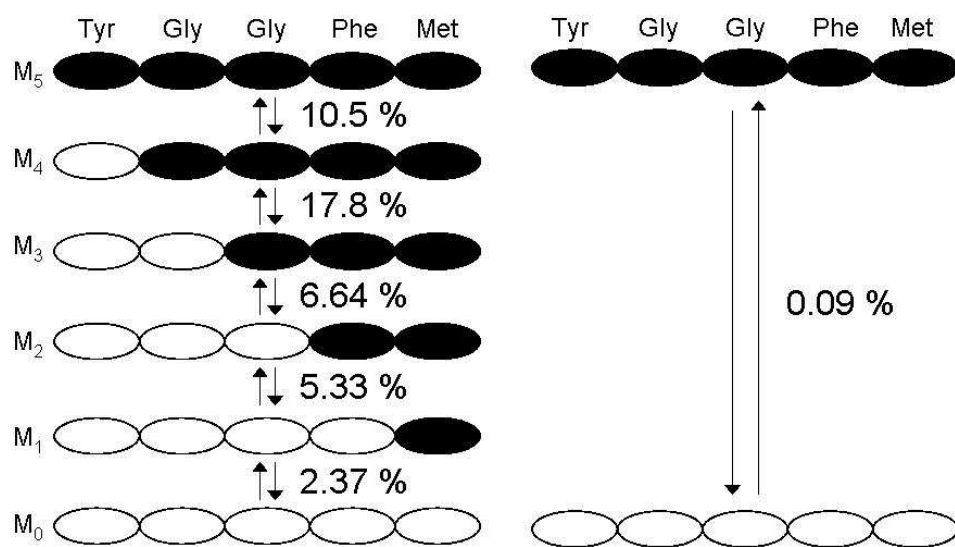


Figure 1: Two different ladders used to exchange all-atom with united-atom met-enkephalin. Residues are depicted with ovals—open corresponds to an all atom representation, filled to united atom. The ratios of successful to attempted exchanges between each level are indicated by the percentages.

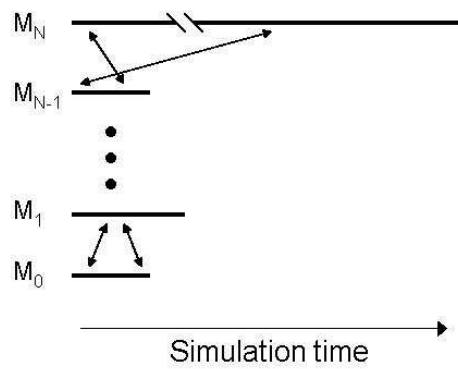


Figure 2: Schematic representation of top-down exchange. Thick horizontal lines are simulation trajectories (labelled “ M_i ” for model “ i ”) and arrows represent exchanges. The M_i may differ in resolution, temperature, or both. Notice that the top level simulation may be considerably longer than the others.




Figure 3: Contours of the potential surface $U(x, y)$, described by Eq. (3). Here we have reduced w to 10 so that both wells are visible in the figure.

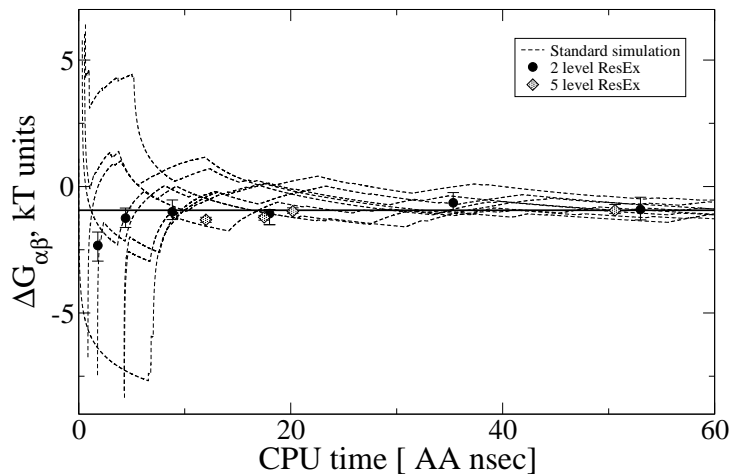


Figure 4: Comparison of different ResEx protocols for dileucine. Plotted are free energy difference estimates between the α and β states as a function of total CPU cost. The dashed lines are individual runs generated by standard Langevin dynamics (no exchange), and the solid horizontal line is the average of 4 independent 150 nsec Langevin dynamics simulations. The solid circles are ResEx results from the two level ladder from Ref.,¹ and the diamonds are the ResEx results from the five level ladder, averaged in each case over 8 independent runs. The error bars give the range of the 8 independent runs. The ResEx data points are displaced from the origin to accurately reflect the time invested in generating the top level and all intermediate level distributions. The exchange schedule for ResEx has not been optimized.

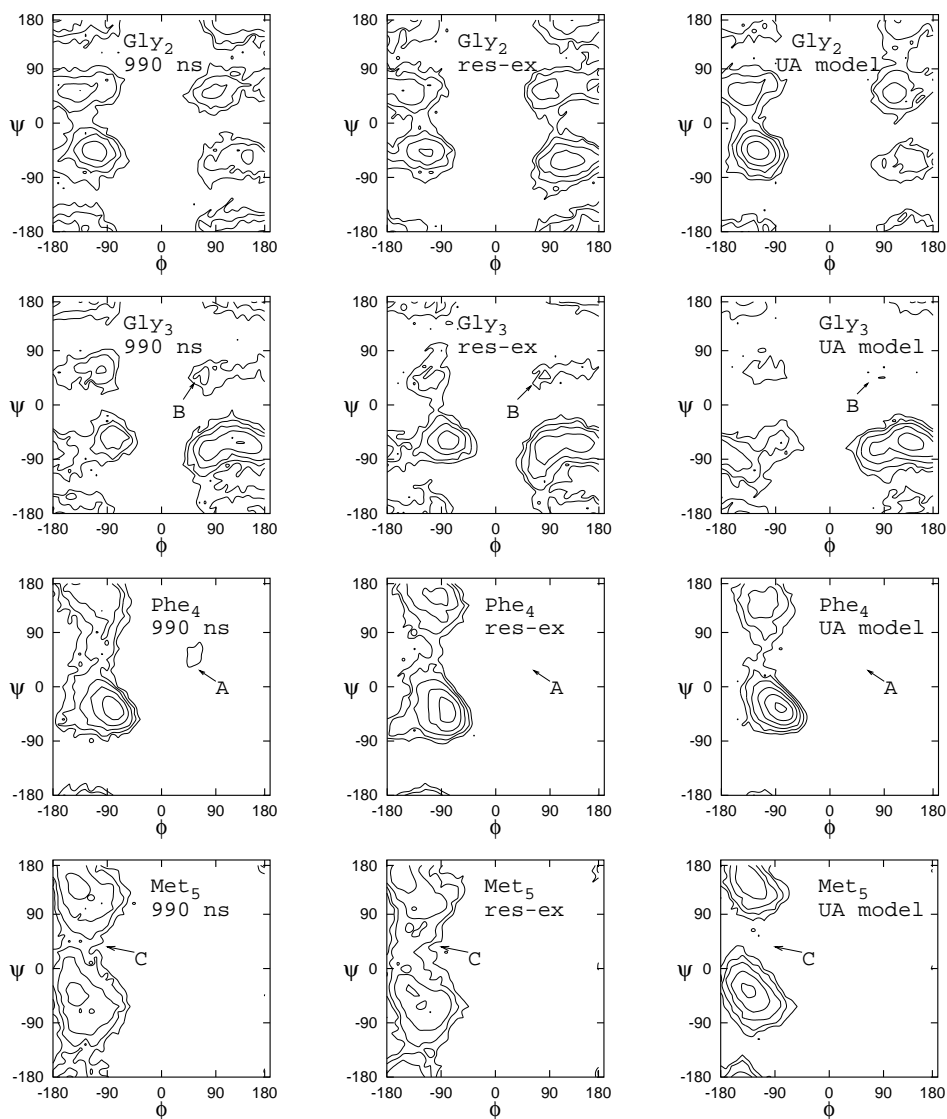


Figure 5: Ramachandran histograms for met-enkephalin. The left column is a 990 nsec Langevin dynamics simulation at all-atom resolution, without resolution exchange; the middle column is the all-atom level (M_0) from resolution exchange as described in the text; the right column is the top-level united-atom simulation (level M_5) used for the resolution exchange simulation shown in the middle column. Note that since the peptide is unblocked, there are only 4 pairs of $\phi - \psi$ dihedrals. Res-ex fails to “find” one region (labelled “A”) not present in the top-level simulation, but finds two others (labelled “B” and “C”).




Figure 6: Ladder combining exchange between all-atom and united-atom met-enkephalin with tempering of reduced resolution simulations. Residues are depicted with ovals—open corresponds to an all atom representation, filled to united atom. The ratios of successful to attempted exchanges between each level are indicated by the percentages. The temperature of each level is indicated on the right.