

Linear-Scaling Soft-Core Scheme for Alchemical Free Energy Calculations

Floris P. Buelens^[a] and Helmut Grubmüller^[a]

Alchemical free energy calculations involving the removal or insertion of atoms into condensed phase systems generally make use of soft-core scaling of nonbonded interactions, designed to circumvent numerical instabilities that arise from weakly interacting “hard” atoms in close proximity. Current methods model soft-core atoms by introducing a nonlinear dependence between the shape of the interaction potential and the strength of the interaction. In this article, we propose a soft-core method that avoids introducing

such a nonlinear dependence, through the application of a smooth flattening of the potential energy only in a region that is energetically accessible under normal conditions. We discuss the benefits that this entails and explore a selection of applications, including enhanced methods for the estimation of free energy differences and for the automated optimization of the placement of intermediate states in multistage alchemical calculations. © 2011 Wiley Periodicals, Inc. *J Comput Chem* 33: 25–33, 2012

Keywords: free energy calculation • molecular dynamics • soft-core scaling • thermodynamics • algorithms

Introduction

The calculation of free energy differences from molecular simulation is of interest to a variety of fields, including computational biophysics and drug design. In these fields, ongoing advances in computational power and methodology have at least partially addressed the shortcomings that became apparent after a first wave of interest starting in the late 1980s,^[1] to the extent that free energy calculation methods, despite difficulties in implementation that continue to hinder their routine application, are developing into tools regularly used to study aspects of molecular recognition and energetics in complex biomolecular systems.^[2]

One major branch of free energy calculation methodology is that of so-called alchemical methods, where free energy changes are calculated along physically impossible but computationally realizable transformations between related molecules, by the gradual transmutation of atomic properties and interactions. The interactions of a set of atoms are described by two closely related potential functions V_1 and V_2 , each representing a state of interest. In the most straightforward implementation, the system is transformed between states by linear interpolation of the two potentials,^[3,4] invoking a coupling parameter λ :

$$V(\lambda) = \lambda V_1 + (1 - \lambda)V_2. \quad (1)$$

Free energy differences can then be calculated using any of a number of available methods.^[5]

Transformations of interest often involve the addition or removal of groups of atoms, such as amino acid side chains or functional groups of a small molecule. In these transformations, straightforward linear interpolation schemes fail when weakly interacting atoms come into close proximity, as a result of the unbounded nature of the nonbonded (van der Waals and electrostatic) potentials at short range. In current applications, this issue is almost universally remedied with the use of “soft-core” scaling

schemes,^[4,6–8] where the van der Waals and electrostatic potentials are modified such that the full-strength, unscaled potential matches the unmodified interaction, whereas in intermediate states the potential is modified such that overlapping atoms no longer give rise to singularities.

While soft-core potentials solve the overlap problem, one desirable property of linear scaling schemes is lost: when the intermediate potential at any given state between end points is represented by a linear combination of the two end point potential functions, it is trivial to extrapolate the internal energy at any other point in the transformation. Current soft-core potentials introduce a nonlinear dependence on the mixing parameter,^[6,8] and calculation of the internal energy in alternate states requires a re-evaluation of all soft-core interactions, either on the fly during the simulation or by later reprocessing of saved configurations of the system.

We here describe a formulation of a soft-core scaling scheme for nonbonded interactions, referred to as linear soft-core (LSC), which has been implemented in the molecular dynamics (MD) package Gromacs.^[9] LSC is designed to combine the advantages of both approaches, namely the numerical integrity of nonlinear-scaling soft-core potentials and the straightforward extrapolation of internal energy in other states that follows from linear scaling. This will be followed by examples of potential applications, including enhancements to methods for the estimation of free energy differences and a scheme to obtain uniform ensemble overlap between neighboring intermediate states.

[a] F. P. Buelens, H. Grubmüller
Department of Theoretical and Computational Biophysics,
Max Planck Institute for Biophysical Chemistry, Am Fassberg 11,
37077 Göttingen, Germany
E-mail: fbuelen@gwdg.de

Concept

The basis of this work is a modification to nonbonded potential functions designed to fulfil the following two conditions.

First, in the initial and final thermodynamic states [“end states,” e.g., $\lambda = 1$ for V_1 and $\lambda = 0$ for V_2 ; eq. (1)], the configurational ensembles generated must be indistinguishable from those generated by the respective unmodified (“native”) potentials. If this condition is fulfilled, the free energy difference associated with a hypothetical switch from the modified to the native potential in the end state is by definition zero, and no correction is needed.

Second, those portions of the nonbonded potentials (V_{nb}) that are not energetically accessible under standard conditions, and thus have no influence on the native ensemble, should be optimized for good numerical behavior when scaled linearly. To this end, the linear soft-core potential V_{LSC} is divided into two regions, native V_{nb} and “capped” V_{cap} , separated at a threshold distance r_{cap} , where

$$\begin{aligned} V_{\text{LSC}}(r) &= V_{\text{nb}}(r) & \text{for } r \geq r_{\text{cap}} \\ V_{\text{LSC}}(r) &= V_{\text{cap}}(r) & \text{for } r < r_{\text{cap}}. \end{aligned}$$

V_{nb} can refer to either the van der Waals or the electrostatic portion of the nonbonded potential, or to both. When both are considered simultaneously, a single value of r_{cap} is applied to all nonbonded interactions. The following discussion assumes that both components are scaled simultaneously, under the control of a single parameter λ .

The cap potential V_{cap} is designed to match the native potential at r_{cap} , with the region below r_{cap} altered to eliminate short-range singularities. Crucially, r_{cap} is chosen so that the capped

region of the potential is inaccessible at standard temperatures and pressures, leaving the accessible regions of the nonbonded potential unmodified.

Defining r_{cap}

The first step is to define, for each pairwise soft-core interaction, the value of r_{cap} , which sets a threshold inside which a full, unscaled nonbonded interaction is judged to be energetically inaccessible at standard temperature and pressure. Considering that these rules will be evaluated a large number of times in the inner nonbonded loop during MD integration, excessive complexity should be avoided.

A first inaccessible region of configurational space is delineated by large positive van der Waals energies. We postulate that interactions with a van der Waals energy higher than a chosen threshold V_{max} will be inaccessible at standard temperature and pressure (Fig. 1). Given a sufficiently conservative choice of V_{max} , the region of the nonbonded potential where $V_{\text{vdw}} > V_{\text{max}}$ may be altered without affecting the accessible energy landscape of the end state.

This first definition will not cover interactions involving hydrogen or other atoms with weak or no van der Waals terms. For these atoms, which are normally shielded by the parent atoms to which they are bonded, the removal of shielding interactions will expose either an electrostatic singularity (Fig. 2d) or a van der Waals radius narrower than the radius actually accessible under standard conditions (Fig. 2g). Such exceptions are caught by defining a minimum allowed separation r_{min} ; we determine that no full-strength pairwise nonbonded interaction at standard temperature and pressure would involve distances shorter than r_{min} , and cap the potential below r_{min} accordingly.

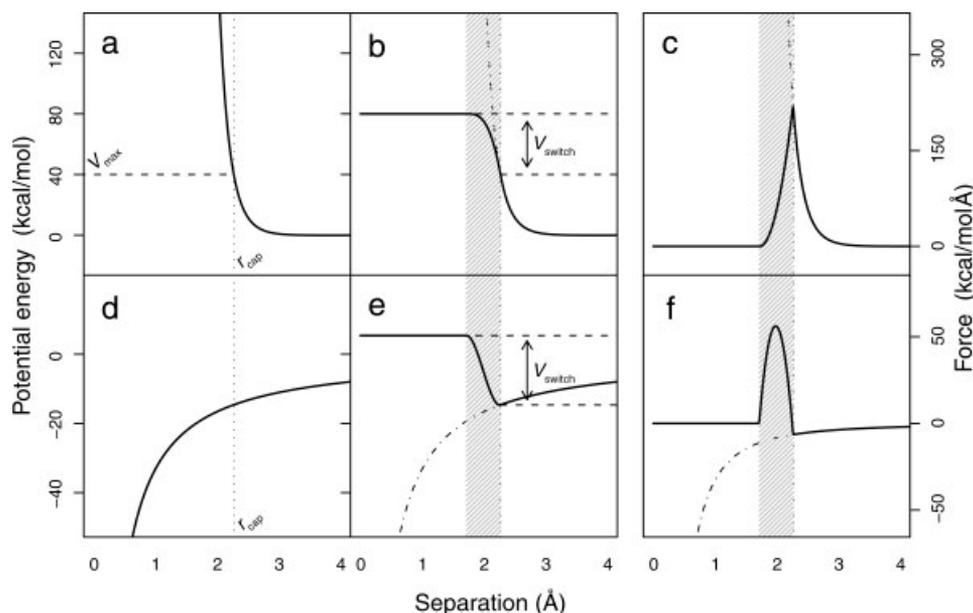


Figure 1. LSC applied to Lennard-Jones (a–c) and Coulomb (d–f) functions, with $V_{\text{max}} = 40$ kcal/mol, for a sample atom pair where charge product $q_i q_j = -0.1$, and Lennard-Jones parameters $\sigma_{ij} = 3.75$ Å and $\epsilon_{ij} = 0.086$ kcal/mol. a and d) Unmodified potential energy as a function of separation $r \times r_{\text{cap}}$ (dotted line) is set as the distance where V intersects a chosen threshold Lennard-Jones energy V_{max} . The derived value of r_{cap} is used to modify both the Lennard-Jones and Coulomb potentials. b and e) Modified (LSC) potential energy, with a dot-dash line showing the original potential. The shaded area, of width r_{switch} and bounded to the right by r_{cap} , represents the range over which the potential is flattened. c and f) Force for LSC (solid line) and unmodified (dot-dash line) potentials.

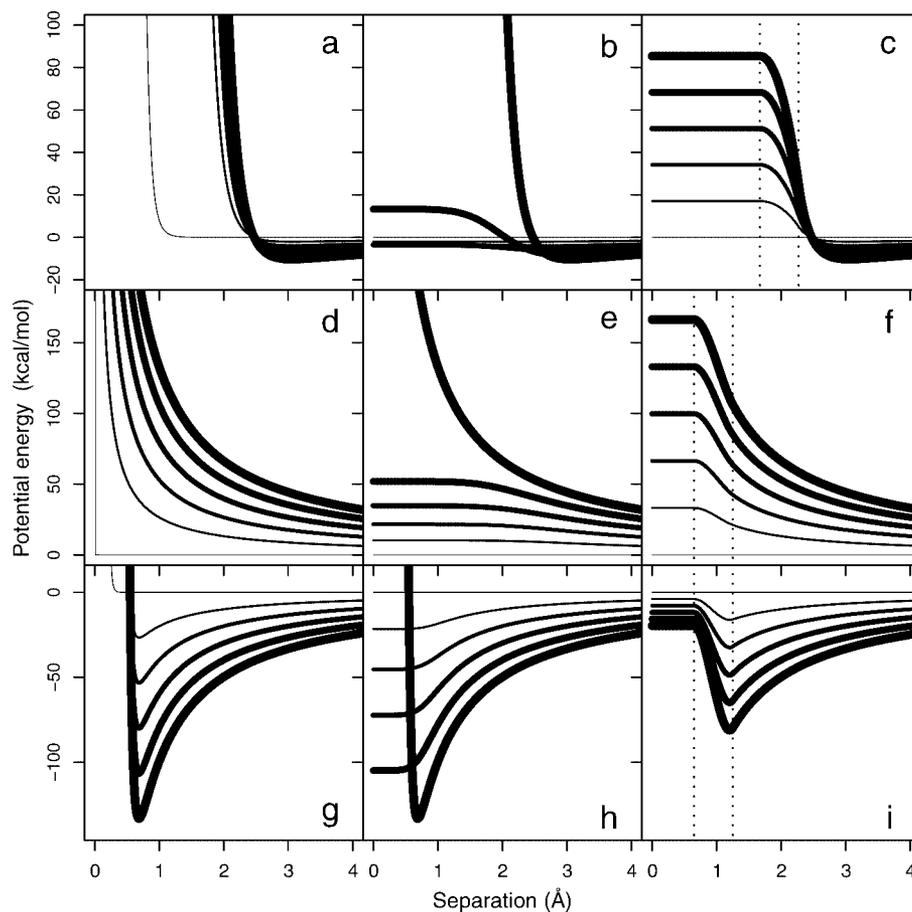


Figure 2. Unmodified, nonlinear soft-core and LSC nonbonded potential energy (Lennard-Jones + Coulomb) for three sample atom pairs. Scaled by values of $\lambda = 1, 0.8, 0.6, 0.4, 0.2,$ and 10^{-5} , where $\lambda = 1$ represents the full-strength, unscaled potential. Heavier lines indicate higher values of λ . a, d, and g) Simple linear scaling of Lennard-Jones and Coulomb potentials. b, e, and h) Nonlinear soft-core potential of Beutler et al.^[6] ($\alpha = 0.5$). c, f, and i) LSC potential, with r_{cap} and r_{switch} indicated with dotted line; $V_{\text{max}} = 40$ kcal/mol, $r_{\text{min}} = 1.25$ Å. a–c) Sample carbon–carbon interaction with $q_i q_j = -0.1$; interaction capped below 2.704 Å. d–f) Atom pair interaction with no van der Waals term, $q_i q_j = 0.4$; interaction capped below r_{min} . g–i) General amber force field (GAFF) hn–hn interaction with $q_i q_j = -0.3$. Weak Lennard-Jones term; interaction capped below r_{min} .

The above criteria allow r_{cap} to be defined for any arbitrary nonbonded potential function. While r_{cap} , through the threshold energy V_{max} , is defined with reference to the van der Waals energy for an atom pair, the derived value can be applied to any nonbonded interaction, as shown with the example of the Coulomb potential (Fig. 1).

Cap potential

The cap potential may, in principle, take any form which eliminates the singularity at zero separation, although the potential energy and its derivative (the force) must be continuous with respect to separation distance to be usable in MD simulation.

A straightforward approach could be to bring the force linearly to zero below r_{cap} , resulting in a quadratically smooth leveling of the potential energy curve. Such an approach would result in a manageable cap on the nonbonded potential in many cases.

However, this approach does not take into account that the nonbonded force at r_{cap} need not be repulsive. In the example of an alchemically scaled interaction between a negatively charged atom and a hydrogen atom with no Lennard-Jones term, the

force at r_{cap} will be strongly attractive. In this case, a simple force interpolation to zero will result in a “negative cap,” where the energy for superimposed atoms is more favorable than that at the capping distance, which risks violating the principle that the capped portion of the potential must be inaccessible under standard conditions.

To ensure a well-defined capped potential surrounding each atom, such that the potential energy of the capped portion of the LSC potential is guaranteed to be less favorable than that at r_{cap} , cubic polynomial functions are defined for both the van der Waals and electrostatic potentials with the following properties:

$$V_{\text{LSC}}(r) = \begin{cases} V(r_{\text{cap}}) + V_{\text{switch}} & \text{for } r \leq (r_{\text{cap}} - r_{\text{switch}}) \\ ar^3 + br^2 + cr + d & \text{for } r_{\text{cap}} - r_{\text{switch}} < r < r_{\text{cap}} \\ V(r) & \text{for } r \geq r_{\text{cap}} \end{cases}$$

$$f_{\text{LSC}}(r) = \begin{cases} 0 & \text{for } r \leq (r_{\text{cap}} - r_{\text{switch}}) \\ -(3ar^2 + 2br + c) & \text{for } r_{\text{switch}} < r < r_{\text{cap}} \\ f(r) & \text{for } r \geq r_{\text{cap}} \end{cases}$$

where $V(r)$ and $f(r)$ represent the unmodified potential energy and force functions, $V_{\text{LSC}}(r)$ and $f_{\text{LSC}}(r)$ are the respective linear-scaling soft-core equivalents, V_{switch} is the final energy penalty

over and above the potential energy at r_{cap} for a separation of zero, and r_{switch} is the distance over which V_{switch} is reached (Fig. 1). Parameters a – d are calculated to match the boundary conditions at r_{switch} and r_{cap} .

Computational Details

The decision whether to cap each nonbonded interaction must be taken on a pair-by-pair basis. While evaluation of conditional statements in the inner loop of the code will generally impair performance, the fact that in Gromacs, alchemically scaled interactions are channeled into a separate pair list and evaluated using a distinct set of routines means that only a small minority of force evaluations is affected.

For each atom pair interaction:

- The native function (electrostatic and van der Waals) for the atom pair is evaluated.
- If the van der Waals potential energy for the pair exceeds V_{max} , or the separation between the atom pair is less than r_{min} , r_{cap} is set so that $V(r_{\text{cap}}) = V_{\text{max}}$ or $r_{\text{cap}} = r_{\text{min}}$. Both the native electrostatic and van der Waals potentials are discarded and replaced by respective capped versions.
- Energy and force, native or capped, are then scaled linearly according to the appropriate value of λ .

At this stage, a check is performed as to whether the capped potential has been applied to a full strength, end state interaction (e.g., a capped interaction was multiplied by a value of $\lambda = 1$). This circumstance represents a violation of the principle that the capped portion of the LSC potential must be inaccessible for the full strength potential, and a warning is issued. If no such warnings are received over the course of a simulation at $\lambda = 0$ or $\lambda = 1$, this condition is met and the integrity of the end state is guaranteed.

Choice of parameters

The LSC scheme includes the adjustable parameters V_{max} , r_{min} , V_{switch} , and r_{switch} , which, respectively, determine the distance at which regular interactions are capped, the minimum allowed distance before a cap is applied, the total energy of the smooth cap potential, and the distance over which the cap is applied. These parameters are tuneable to meet the needs of the application. We here discuss parameter choice for a typical system (condensed phase, all-atom force field at 298 K and atmospheric pressure).

As discussed above, the choices of V_{max} and r_{min} are dictated by the need to define r_{cap} such that the altered region of the unscaled potential is energetically inaccessible. A intuitive starting point for V_{max} would be to consider the collision of simple Lennard-Jones particles and to choose an energy level unlikely ever to be observed on MD time scales. A value of 15 kcal/mol would fulfil this criterion, as less than one event per 10^{11} would be expected to exceed this energy. However, when considering charged particles, dipoles and solvent shielding in the condensed phase, atom pair interaction energies will routinely exceed this value. Considering these issues, we found a value of $V_{\text{max}} = 40$ kcal/mol, taking into consideration only the van der Waals component of the interaction energy, was sufficiently

conservative for a contrived test simulation setup consisting of a mixture of water and high concentrations of monovalent and divalent ions and polar and nonpolar organic molecules, with no violations of the rule observed in 20 ns of simulation. The value $r_{\text{min}} = 1.25$ Å was empirically derived from the same test system. Less conservative values can also be successfully used.

While the quantity V_{switch} on the one hand serves to ensure that overlapping states are always unfavorable relative to the native potential, the choice of V_{switch} and r_{switch} is also dictated by the need to bring the force to zero in a controlled manner. Excessively small values will lead to violent changes in the force, culminating in the limit of $V_{\text{switch}} \rightarrow 0$ or $r_{\text{switch}} \rightarrow 0$ in a discontinuity, resulting in integration errors and possible simulation instability. The suggested values of $V_{\text{switch}} = 40$ kcal/mol for the van der Waals and 20 kcal/mol for the electrostatic potential, and $r_{\text{switch}} = 0.4$ Å were chosen conservatively to avoid changes in the force any steeper than those imposed by the original potential.

Applications

Assuming linear scaling of other components of the potential function, like bonded interactions, the LSC potential permits the hybrid potential to be composed entirely of linear combinations of two end point potential functions, without the loss of numerical integrity this normally implies. As a consequence, it is trivial to extrapolate the internal energy of any arbitrary hybrid state of the system using only the derivative $dV/d\lambda$.

Such extrapolation of the internal energy for arbitrary states along a linear transformation pathway is useful for a number of techniques which make use of ensemble reweighting or otherwise require knowledge of the energy difference associated with a hypothetical switch to another state. The remainder of this section will explore examples that seek to illustrate potential applications and advantages over non-LSC scaling schemes.

Free energy calculation techniques

Broadly, three categories of methods are commonly used for the calculation of free energies from alchemical simulations at equilibrium: thermodynamic integration (TI),^[10] exponential averaging (EXP) with the Zwanzig formula^[11] and Bennett's acceptance ratio method (BAR),^[12] along with a recently developed elaboration, multistate BAR or MBAR.^[13] All three methods can benefit from the use of the LSC potential.

EXP and BAR. Both EXP and BAR depend on the calculation of instantaneous work values between neighboring states of a multistage alchemical free energy calculation. For each pair of states A and B, instantaneous work values $W_{A \rightarrow B} = E_B - E_A$ are calculated for configurations generated in ensemble A, and values $W_{B \rightarrow A} = E_A - E_B$ for configurations generated in ensemble B. On the basis of these values, free energy differences may be estimated, either in a directional manner from A to B and *vice versa* with EXP, or by means of statistically optimized overlap sampling with BAR. Non-linear soft-core methods require such work values to be calculated at the level of the inner loop of

the nonbonded interaction calculation, either by postprocessing of stored trajectories with a different potential function or at run-time. This need not represent a significant burden in terms of computational cost or code complexity, although linear extrapolation would reduce both to some extent.

In contrast, the practicality of MBAR is greatly enhanced by LSC scaling. MBAR analysis requires the availability of instantaneous work values associated with every possible pairing of states in the transformation; for n intermediate states, this implies the calculation of $n - 1$ work values at regular intervals in each simulation. Although the associated computational cost for run-time calculation need not be large (only those interactions that are alchemically scaled need to be re-evaluated, and work values need not be calculated at every step), there is a significant overhead in code and workflow complexity, and to some extent for long simulations with large values of n , in terms of data storage and memory requirements. Linear scaling avoids these overheads; only $dV/d\lambda$ must be calculated and stored, and instantaneous work values between pairs of states can be calculated straightforwardly during MBAR analysis.

Thermodynamic integration TI is another widely applied free energy calculation technique; the free energy for an alchemical transformation may be estimated as the integral of $dV/d\lambda$ along the transformation pathway. All implementations of TI are to some extent subject to bias arising from the numerical estimation of the integral based on data from a finite number of intermediate states. Basic schemes like trapezoidal integration are particularly weak in this respect; higher-order interpolation schemes have been shown to perform better.^[14] Here, we explore a higher-order spline-based interpolation scheme, which makes use of ensemble reweighting techniques made possible by the linear-scaling nature of the LSC potential.

Fundamentally, the implementation of a TI scheme requires only the calculation of the derivative of the potential function with respect to λ . With linear mixing as per $V_\lambda = (1 - \lambda)V_{\text{start}} + \lambda V_{\text{end}}$, the derivative is simply $V_{\text{end}} - V_{\text{start}}$; these values of $dV/d\lambda$ are collected during the simulation and stored for later analysis.

The average over a thermodynamic ensemble B of any thermodynamic observable x can be estimated on the basis of a sample from a closely related ensemble A according to:^[15]

$$\langle x \rangle_B = \frac{\langle x \exp[W_{A \rightarrow B}/k_B T] \rangle_A}{\langle \exp[W_{A \rightarrow B}/k_B T] \rangle_A} \quad (2)$$

This reweighting approach can be applied to estimate the value of $dV/d\lambda$ at any arbitrary point of the transformation. In principle, this approach could be used to construct a fine-grained estimate of the shape of the full integral, although the statistical quality of the estimated values will reduce as work values $W_{A \rightarrow B}$ increase with decreasing overlap between ensembles A and B. We here apply an alternative approach: estimates of $dV/d\lambda$ are made for values of λ only slightly outside the simulated ensembles, and used to estimate higher derivatives of the free energy with respect to λ with a finite difference approach. These higher derivatives can then be used as a physically derived input for higher-order interpolation schemes.

To estimate second and higher derivatives of the free energy with respect to λ by finite difference, we state the fundamental relationship^[10]:

$$G'(\lambda) = \left\langle \frac{dV}{d\lambda} \right\rangle_\lambda, \quad (3)$$

where the angle brackets with subscript λ indicate an ensemble average collected for a given value of λ . Applying eq. (2), we can reweight the ensemble generated at λ to estimate ensemble averages of $dV/d\lambda$ for $\lambda - \Delta\lambda$ and $\lambda + \Delta\lambda$, where $\Delta\lambda$ is chosen as small as practically possible. $G''(\lambda)$ is then approximated as:

$$G''(\lambda) \approx \frac{G'(\lambda + \Delta\lambda) - G'(\lambda - \Delta\lambda)}{2\Delta\lambda}. \quad (4)$$

Iterative application following the pattern:

$$G'''(\lambda) \approx \frac{G''(\lambda + \Delta\lambda) - G''(\lambda - \Delta\lambda)}{2\Delta\lambda} \quad (5)$$

allows the estimation of derivatives to arbitrary order.

Any numerical integration scheme implicitly involves predicting the shape of the curve between sampled data points; we here seek to extrapolate $G' = \langle dG/d\lambda \rangle$ by spline interpolation. Splines S of order m can be used to interpolate G' for each section $\lambda_i \rightarrow \lambda_{i+1}$:

$$G'(\lambda) \approx S(\lambda) = \sum_{j=0}^m c_j (\lambda - \lambda_i)^j, \quad (6)$$

where $c_{0..m}$ are parameters fitted such that for each piecewise polynomial from $\lambda = \lambda_i$ to $\lambda = \lambda_{i+1}$, derivatives of S (e.g., of G') up to order $m - 2$ are continuous between neighboring polynomials. When fitting derivatives of S of order x , a spline of order $m = 2x + 1$ results in a unique solution for each polynomial segment. Note that interpolation incorporating no derivatives of S (e.g., G' and no higher) corresponds to $m = 1$, which reduces to trapezoidal integration.

For a free energy calculation with data sampled in n intermediate states, the TI estimate is then given by

$$\begin{aligned} \Delta G &= \sum_{i=0}^{n-1} \int_{\lambda_i}^{\lambda_{i+1}} G'(\lambda) d\lambda \\ &\approx \sum_{i=0}^{n-1} \int_{\lambda_i}^{\lambda_{i+1}} S(\lambda) d\lambda, \end{aligned} \quad (7)$$

where the integrals of each segment $S(\lambda)$ can be evaluated analytically.

Automated balancing of ensemble overlap

The degree of overlap along the chain of ensembles that describes a multistage transformation is a key determinant of the efficiency of estimation of free energy differences.^[12] By extension, a transformation pathway designed for optimal efficiency should attempt to balance the extent to which neighboring distributions along the chain overlap.^[16,17] We here propose a method

which automates the generation of a chain with balanced ensemble overlap, building on the observation that the probability of replica exchange acceptance between neighboring ensembles is a quantitative measure of the extent of their overlap.^[18]

Replica exchange between intermediate stages in an alchemical transformation is a form of Hamiltonian replica exchange,^[19,20] where the difference in internal energy between related potential functions is used as the exchange quantity, as opposed to the more widely used temperature temperature replica exchange technique.^[21] Hamiltonian replica exchange can allow barriers to sampling along the λ coordinate to be crossed more readily and thereby enhance sampling.

The probability p of exchange move acceptance between two configurations i and j , belonging to two intermediate states A and B with potential functions V_A and V_B , is given by

$$p = \exp\left(\frac{V_A(i) - V_B(i) + V_B(j) - V_A(j)}{k_B T}\right). \quad (8)$$

The use of linear scaling means that $V_B(i)$ and $V_A(j)$ are readily available from only the values of $dV/d\lambda$ with no need for inner loop calculations, avoiding associated computational and code maintenance overheads. Extension to all-pairs exchange^[22] likewise requires no additional computation of energy differences.

To estimate the average exchange probability ($\langle p \rangle$), and thereby the degree of ensemble overlap between two sampled ensembles A and B ,^[18] we repeatedly apply eq. (8) with a random sample of configurations i and j , respectively drawn from A and B . The accumulated average probability will approach the underlying statistical mechanical probability of exchange between configurations of ensembles A and B .

With a linear-scaled potential, this approach can easily be extended to incorporate hypothetical intermediate states, for which no data have been collected directly. Data for such a hypothetical ensemble C , closely related to ensemble B , may be generated by taking a sample from ensemble B , with the probability of picking each configuration weighted according to $\exp(W_{B \rightarrow C}/k_B T)$.^[15,23] The sample from hypothetical ensemble C can then be used to estimate the average exchange frequency between configurations of C and any other ensemble, whether explicitly simulated or hypothetical.

Given an alchemical transformation consisting of n states i , an initial simulation with arbitrarily distributed intermediates will yield data for which the acceptance probability ($\langle p \rangle$) between states i and $i + 1$ can be observed, giving a vector of $n - 1$ successive acceptance probabilities ($\langle p_{i \rightarrow i+1} \rangle$). Considering successive exchanges in λ as a cumulative chain of probabilities, a target acceptance probability p_t can then be defined as the $(n - 1)$ th root of the product:

$$P_t = \sqrt[n-1]{\prod_{i=0}^{i < n-1} \langle p_{i \rightarrow i+1} \rangle}. \quad (9)$$

Leaving the (physically relevant) first and last states intact, the intervening space on the λ axis can then be filled with the remaining $n - 2$ intermediate states, by instating a new value

of λ for each intermediate i and iteratively adjusting it until the probability of exchange with ensemble $i - 1$ matches the target probability. Iterative application of this procedure with successively refined values of p_t and/or fresh data sampled from the new set of intermediate states should generate a set of λ values such that the probability of exchange between ensembles representing neighboring intermediate states uniformly matches the target value p_t .

Results and Discussion

The LSC potential was implemented in Gromacs 4.0.5. To demonstrate functioning of the methods proposed, we considered the decoupling of a small molecule, indole, from a TIP3P water environment. The measurement system consisted of indole, parameterized according to the general amber force field,^[24] in a periodic rhombic dodecahedron with 1011 water molecules. MD simulations were performed using particle mesh Ewald for full electrostatic integration, with a real-space cutoff of 10 Å, interpolation order 4 and a grid spacing of 1.2 Å. Van der Waals interactions were truncated using a switching function acting from 8 to 9 Å; an analytical correction for neglected long-range van der Waals interactions was applied to energy and pressure. Bond lengths involving hydrogen were constrained and a time step of 2 fs was used. The alchemical transformation was performed using 16 intermediate states describing the transformation from $\lambda = 0$ (indole decoupled from surroundings, with intramolecular interactions unmodified) to $\lambda = 1$ (all van der Waals and electrostatic interactions at full strength). Replica exchanges were attempted every 100 steps. No stability problems were encountered in MD based on the LSC potential. LSC parameters of $V_{\max} = 40$ kcal/mol, $r_{\min} = 1.1$ Å and $r_{\text{switch}} = 0.4$ Å were used. V_{switch} was set to 40 kcal/mol for the van der Waals and 20 kcal/mol for the electrostatic potential.

Ensemble overlap

The spacing of intermediate simulations across the λ axis was optimized with the scheme described earlier. Consistency of the resulting distribution was assessed by repeating the protocol five times with different initial distributions; intermediates for the first replicate were uniformly distributed from 0 to 1, and for the remaining four were placed at random. Each of the five replicates of 16 intermediates was initially simulated for 40 ps, and the resulting samples were used to calculate the probability of exchange between the neighboring ensembles and to estimate the λ spacing required for uniform replica exchange probability.

After the first round of sampling, all five replicas showed a negligible ($< 10^{-3}$) probability of exchange between intermediate $i = 0$ and $i = 1$, and probabilities ranging from 0.24 to 0.995 for exchanges in the range $i = 1$ to $i = 15$ (Fig. 3a). The apparent lack of ensemble overlap between the first and second intermediates can be interpreted in the context of the cavity formation process that occurs on coupling of a molecule to surrounding solvent. With the LSC parameters used, this process occurs approximately in the region $0 < \lambda < 0.02$; neither uniform nor random spacing placed any intermediate state in this region.

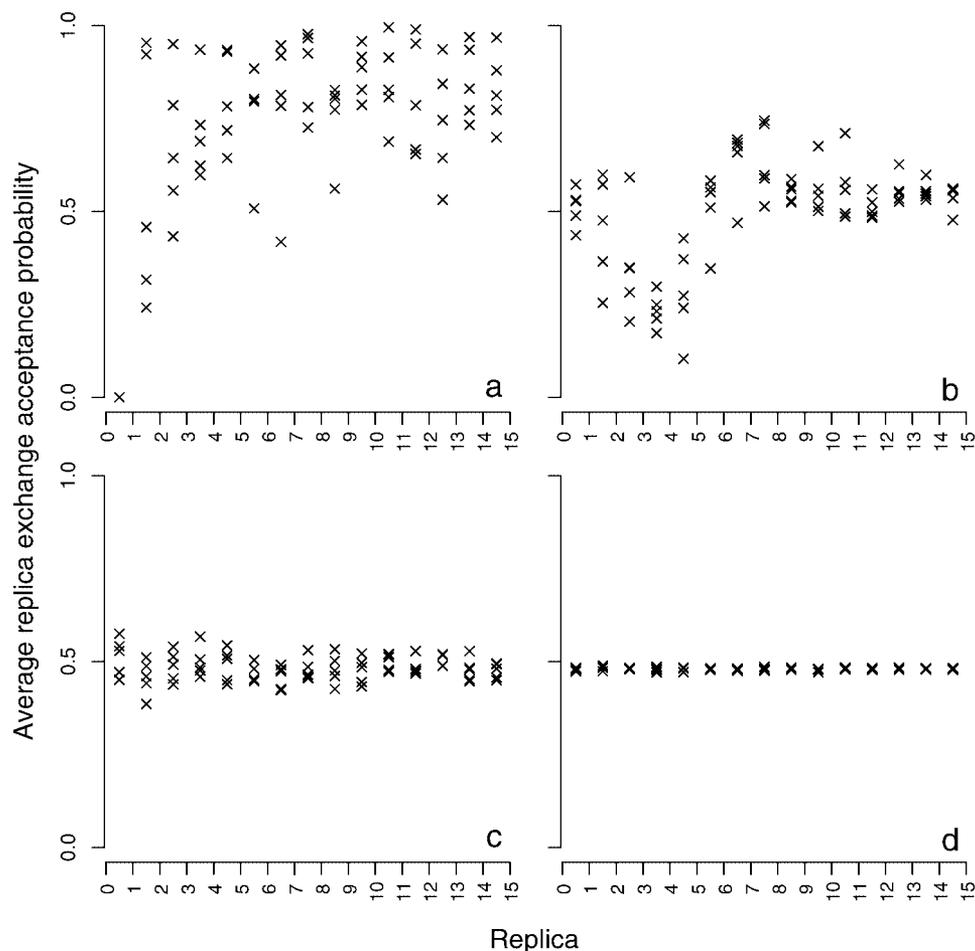


Figure 3. Iterative adjustment of intermediate states showing convergence toward uniform replica exchange acceptance probability. Each mark represents the average acceptance probability for exchange attempts between the ensembles to the left and right, with five replicates in each instance. a) Acceptance probabilities prior to optimization; note negligible probability of exchange between $i = 0$ and $i = 1$ and c) Probabilities after one (b) and two (c) refinement cycles, each based on 16×40 ps of sampling. d) Probabilities after exhaustive sampling.

A refined λ spacing based on the initial suboptimal sample was derived according to the described method. After a second 40 ps round of sampling with refined λ values, a distribution of average replica exchange probabilities emerged with a range of 0.1–0.74. After a second refinement of λ spacing, a third 40 ps sample showed convergence of replica exchange probabilities, with 0.484 ± 0.048 over the 75 exchange pairs (five replicates of 15; Fig. 3c).

Having shown apparent convergence on the basis of 80 ps of sampling and two rounds of refinement for this system, convergence with longer samples was assessed; a correct implementation should show perfect uniformity in the large sample limit. 5 ns of data was collected for each of the five replicates of 16 intermediate states each, using the spacing from the second round of refinement. The spacing was again refined on the basis of this sample, and a further 16×5 ns of data were collected. Observed exchange probabilities approached uniformity across all five replicates, at 0.480 ± 0.003 over 75 exchange pairs (Fig. 3d). λ values for the five replicates were tightly converged, with the range of values representing $0.07 \pm 0.05\%$ of the average in each of the 16 intermediate states, indicating the emergence of a single consensus solution.

The validity of the average replica exchange probability as a measure of the overlap between neighboring distributions was tested by examining the distributions of work values for transformations between neighboring ensembles ($W_{A \rightarrow B}$ and $W_{B \rightarrow A}$; Fig. 4). Across the 75 pairs of neighbors, work value distribution overlap determined from kernel density estimates was well converged at 0.62 ± 0.02 .

Thermodynamic integration

To assess the performance of integration schemes based on higher-order derivatives, an exhaustive calculation of the hydration free energy of indole was performed. Ninety-six intermediate states (λ_i) from $\lambda_0 = 0$ to $\lambda_{95} = 1$, distributed for uniform replica exchange acceptance probability according to the method described above, were simulated for 5 ns each. Subsets of the 96-point sample were analyzed by spline interpolation using derivatives of G up to order 5, corresponding to splines up to order 9, and additionally using BAR. Spline-interpolated curves based on subsets of the 96 points qualitatively reproduced the shape of the G' curve better with increasing interpolation order (Fig. 5). The resulting free energy estimates (Table 1)

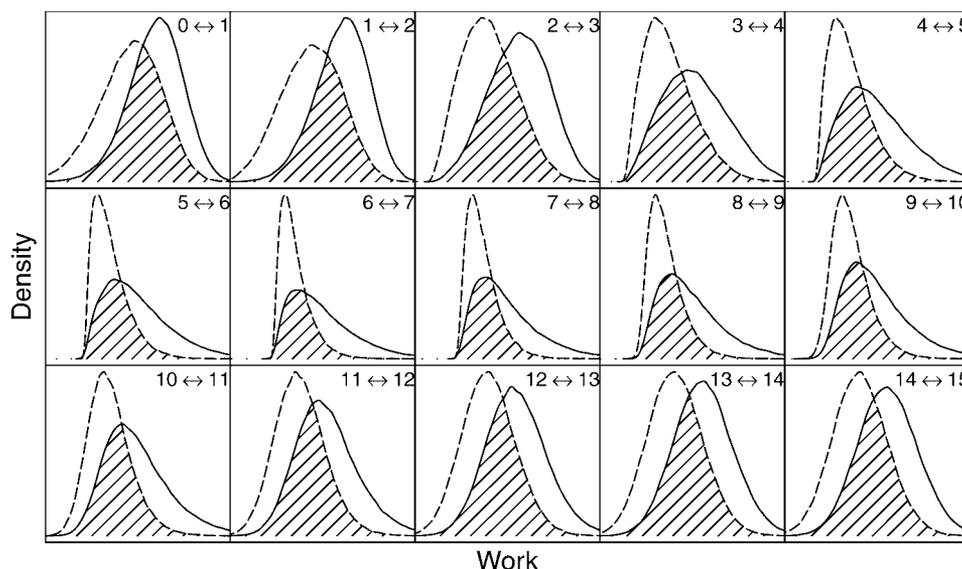


Figure 4. Distributions of instantaneous work values for transformations $A \rightarrow B$ (solid) and $B \rightarrow A$ (broken line) between neighbors after optimization of the 16-state transformation described in the text, with overlap highlighted. Scale omitted for clarity.

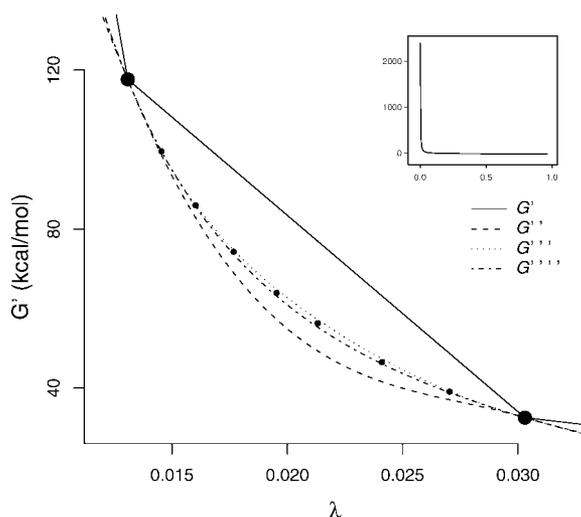


Figure 5. Interpolation of G' showing trapezoidal and progressively higher-order spline schemes. Inset: full TI curve for the hydration of indole. Main plot: detail showing region of greatest curvature. Large circular markers show data points from which splines were derived; intervening small circular markers represent sampled data not included in the analysis. Four curves show the interpolation of G' using progressively higher derivatives of G up to order 4 (G''''), with splines of order $m = 1$ (trapezoidal), 3, 5, and 7. The curve for order 5 ($m = 9$) closely matches that for order 4 and is not shown.

quantitatively show faster convergence as a function of the number of data points provided. Whereas the trapezoidal estimate continues to decrease monotonically even approaching the full set of 96 points, higher-order schemes approached a consensus value using data from as few as nine (order 5) to 13 (order 2) intermediate simulations, representing 7–10 times fewer data points.

Based on these measurements, higher-order spline schemes, incorporating higher derivatives of G derived from ensemble reweighting techniques rendered practical with use of the LSC potential, appear capable of significantly improving the quality

Table 1. Indole hydration free energy (kcal/mol) based on progressively larger subsets of a 96-point sample (5 ns/point), calculated using BAR and spline-interpolated integration incorporating derivatives of G of progressively higher order.

Nr. intermediates	BAR	TI				
		G'	G''	G'''	G''''	G'''''
3	-3.32	40.72	-111.63	826.7	-6968	61696
5	-4.14	4.49	-15.56	13.19	-33.30	56.45
7	-4.06	-0.20	-6.32	-2.46	-5.17	-3.24
9	-4.06	-1.97	-4.73	-3.73	-4.20	-3.98
13	-4.05	-3.13	-4.17	-4.02	-4.05	-4.04
17	-4.05	-3.53	-4.08	-4.04	-4.04	-4.04
25	-4.03	-3.80	-4.04	-4.03	-4.03	-4.03
33	-4.03	-3.90	-4.03	-4.03	-4.03	-4.03
49	-4.03	-3.97	-4.03	-4.03	-4.03	-4.06
96	-4.03	-4.01	-4.03	-4.03	-4.03	-4.06

of TI estimates of free energy differences. However, in this test BAR stands out as the most efficient estimator, approaching the consensus value with as few as five intermediate states.

Comparison with nonlinear soft-core scheme

The LSC scheme described is proposed as an alternative to a number of nonlinear-scaling soft-core interaction schemes, which are in common use. We compared the performance of LSC against the soft-core potential of Beutler et al.,^[6] as implemented in Gromacs. Using the methodology described earlier, but with iterative and time-consuming MD sampling of candidate ensembles in the nonlinear case as opposed to ensemble reweighting in the LSC case, we derived chains of ensembles with uniform overlap, describing the decoupling of indole from an aqueous environment. For this comparison, only the van der Waals component of the insertion process was considered. In both cases, 14 intermediate states were required to produce a chain with replica exchange probability around 0.5; LSC appeared to perform marginally better, with an average exchange frequency of 0.509 with 14 states,

compared to 0.495 for the Beutler et al. potential. We conclude that, at least in this simple example, the phase space overlap properties of the LSC potential are no worse than those for the potential of Beutler et al.

However, it should be noted that while the distribution of intermediate states from $\lambda = 0$ to 1 is fairly uniform for the Beutler et al. potential, λ values with the LSC potential were heavily skewed toward the insertion point (0, 0.004, 0.008, 0.011, 0.014, 0.019, 0.026, 0.037, 0.060, 0.104, 0.194, 0.357, 0.622, and 1, where 0 is the decoupled state). This property may prove disadvantageous in certain nonequilibrium applications where λ is changed between 0 and 1 at a constant velocity.

Conclusions

The LSC potential is designed to combine the numerical integrity of non-LSC potentials at low separations with the practical advantages of linear scaling. By neutralizing the problematic short-range properties of the van der Waals and electrostatic potentials without invoking a nonlinear dependence on the coupling parameter, LSC permits the straightforward application of analysis schemes which require the calculation of instantaneous work values between states, which makes use of ensemble reweighting, or which otherwise require the evaluation of hypothetical alternative potential functions. The examples of LSC-based schemes in this article are intended to explore the flexibility that accompanies linear scaling; it is anticipated that LSC, as part of the wider alchemical free energy calculation tool set, can prove useful for applications beyond those suggested here.

The LSC potential is not dependent on any particular force field or nonbonded potential function. The precise form of the potential is dependent on the tuneable parameters V_{\max} and r_{\min} ; an insufficiently conservative choice of either will compromise the native form of the full-strength potential at the end states, but this outcome will easily be detected and remedied. The real-space component of the particle mesh Ewald electrostatic potential is compatible with LSC.

The scheme described to delineate those areas of phase space that are accessible and inaccessible under standard conditions, as a function of parameters V_{\max} and r_{\min} , is a simplistic one. There is likely to be scope for improvement and tuning of algorithms, either on a pair-by-pair basis or incorporating higher-order information, to sharpen the delineation and thereby improve overlap properties.

A method was described to determine the spacing of alchemical intermediate states that gives rise to uniform ensemble overlap between neighboring ensembles, thereby helping to distribute simulation time optimally across the transformation. To our knowledge, no systematic method has previously been described to automate this process; whereas clearly defined schemes exist to establish which distribution of temperatures is required for a given acceptance probability in the case where temperature is the exchange quantity,^[18] the spacing of intermediate states for Hamiltonian replica exchange with alchemical transformations in particular, and for alchemical free energy calculation in general, remains a process of trial and error. In

published reports of alchemical free energy calculations, the spacing of intermediate states is often presented without detailed supporting rationale; in these cases it can be assumed that the spacing was chosen either more or less by hand (e.g., from the shape of a TI curve) or iteratively with input from sampled data or other statistical measures. The scheme described here, which would be impractical and computationally expensive to implement in combination with non-LSC potentials but which integrates seamlessly with linear-scaling approaches, provides a solution grounded in statistical mechanics, and is suitable for implementation in MD simulation at run-time. Minimizing the need for manual input at this stage represents a step toward overcoming the barriers that continue to preclude alchemical free energy calculations from routine or "black box" use.

Acknowledgments

Funding provided by the BMBF (Grant Number 03154501) is gratefully acknowledged.

- [1] P. A. Kollman, *Chem Rev* 1993, 93, 2395.
- [2] A. Aleksandrov, D. Thompson, T. Simonson, *J Mol Recognit* 2010, 23, 117.
- [3] C. Chipot, A. Pohorille, *Free Energy Calculations*, Springer Series in Chemical Physics; Springer: Berlin, 2007.
- [4] T. Steinbrecher, D. L. Mobley, D. A. Case, *J Chem Phys* 2007, 127, 214108.
- [5] C. D. Christ, A. E. Mark, W. F. van Gunsteren, *J Comput Chem* 2010, 31, 1569.
- [6] T. C. Beutler, A. E. Mark, R. C. Van Schaik, P. R. Gerber, W. F. Van Gunsteren, *Chem Phys Lett* 1994, 222, 529.
- [7] J. W. Pitera, W. F. van Gunsteren, *Mol Simul* 2002, 28, 45.
- [8] M. Zacharias, T. Straatsma, J. McCammon, *J Chem Phys* 1994, 100, 9025.
- [9] B. Hess, C. Kutzner, D. van der Spoel, E. Lindahl, *J Chem Theory Comput* 2008, 4, 435.
- [10] J. G. Kirkwood, *J Chem Phys* 1935, 3, 300.
- [11] R. W. Zwanzig, *J Chem Phys* 1954, 22, 1420.
- [12] C. H. Bennett, *J Comput Phys* 1976, 22, 245.
- [13] M. R. Shirts, J. D. Chodera, *J Chem Phys* 2008, 129, 124105.
- [14] C. Shyu, F. M. Ytreberg, *J Comput Chem* 2011, 32, 134.
- [15] G. M. Torrie, J. P. Valleau, *J Comput Phys* 1977, 23, 187.
- [16] A. Pohorille, C. Jarzynski, C. Chipot, *J Phys Chem B* 2010, 114, 10235.
- [17] D. Wu, D. A. Kofke, *J Chem Phys* 2005, 123, 084109.
- [18] D. A. Kofke, *J Chem Phys* 2002, 117, 6911.
- [19] H. Fukunishi, O. Watanabe, S. Takada, *J Chem Phys* 2002, 116, 9058.
- [20] Y. Sugita, A. Kitao, Y. Okamoto, *J Chem Phys* 2000, 113, 6042.
- [21] Y. Sugita, Y. Okamoto, *Chem Phys Lett* 1999, 314, 141.
- [22] P. Brenner, C. R. Sweet, D. VonHandorf, J. A. Izaguirre, *J Chem Phys* 2007, 126, 074103.
- [23] R Development Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing: Vienna, Austria, 2010; ISBN 3-900051-07-0.
- [24] J. Wang, R. M. Wolf, J. W. Caldwell, P. A. Kollman, D. A. Case, *J Comput Chem* 2004, 25, 1157.

Received: 16 January 2011

Revised: 22 July 2011

Accepted: 11 August 2011

Published online on 26 September 2011