



Exploiting the potential energy landscape to sample free energy

Andrew J. Ballard,¹ Stefano Martiniani,¹ Jacob D. Stevenson,¹ Sandeep Somani^{1,2} and David J. Wales^{1*}

We review a number of recently developed strategies for enhanced sampling of complex systems based on knowledge of the potential energy landscape. We describe four approaches, replica exchange, Kirkwood sampling, superposition-enhanced nested sampling, and basin sampling, and show how each of them can exploit information for low-lying potential energy minima obtained using basin-hopping global optimization. Characterizing these minima is generally much faster than equilibrium thermodynamic sampling, because large steps in configuration space between local minima can be used without concern for maintaining detailed balance. © 2015 John Wiley & Sons, Ltd.

How to cite this article:

WIREs Comput Mol Sci 2015, 5:273–289. doi: 10.1002/wcms.1217

INTRODUCTION

In this contribution, we provide an overview of recent methods that all have their basis in enhancing the sampling of global thermodynamics using knowledge of the underlying potential energy landscape. Making these connections seems particularly timely, as a number of potentially important advances have been made over the last few years, and they share some common characteristics. In particular, the use of low-lying local minima obtained from methods based on global optimisation (Section *Basin-Hopping Global Optimisation*), and the superposition approach (Section *The Superposition Approach*), provide some common themes. Of course, many other methods have been proposed for enhanced sampling,^{1–13} and these will not be reviewed here. The procedures we describe, with foundations in potential energy landscape theory^{14–16} and geometry optimisation, are largely complementary to approaches based on more conventional molecular dynamics and Monte Carlo (MC) schemes.¹⁷

Future work will require benchmarking comparisons of the most promising methods, to guide applications in different fields. For the present purposes, the main aim is to show what has been achieved in recent work based on the potential energy landscape framework, and explain how these new tools are connected.

Most of the tests we have conducted use atomic clusters of N atoms bound by the Lennard-Jones (LJ) potential¹⁸

$$V = 4\epsilon \sum_{i < j} \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^6 \right], \quad (1)$$

where ϵ and $2^{1/6}\sigma$ are the pair equilibrium well depth and separation, respectively. These systems will be denoted LJ_N , and there are a number of sizes where low temperature solid–solid phase transitions have been identified in previous work.^{7–8,14,19–26} In each case, there are competing low-energy morphologies separated by potential energy barriers that are large compared with $k_B T$ at the transition temperature. The relatively simple potential makes these systems ideal benchmarks for analysis of broken ergodicity issues, as well as for global optimisation^{14,19,27} and rare event dynamics.^{28–31} In principle, the same improvements would be obtained for treatments that explicitly consider electronic structure. However, the much greater

*Correspondence to: dw34@cam.ac.uk

¹Department of Chemistry, University Chemical Labs, Cambridge, UK

²Discovery Sciences, Janssen Research & Development LLC, Spring House, PA, USA

Conflict of interest: The authors have declared no conflicts of interest for this article.

computational expense of computing energies and gradients that would be necessary mean that these opportunities have yet to be exploited.

We begin by introducing the superposition approach (Section *The Superposition Approach*), a key component in the methods under review, which allows us to link thermodynamics to the basins of attraction of local minima on the underlying energy landscape. With this background in place, we review in Section *Basin-Hopping Global Optimisation* the basin-hopping (BH) strategy for global optimisation. We then consider strategies for thermodynamic estimation, including replica exchange (REX) (Section *Replica-Exchange Approaches*), Kirkwood sampling (Section *Kirkwood Sampling*), nested sampling (NS) (Section *Superposition-Enhanced Nested Sampling*), and basin-sampling (Section *Basin-Sampling*). In these sections, we emphasize the utility of the superposition approach in enhancing these simulations, and provide evidence for gains in efficiency and convergence from simulation results for benchmark systems.

THE SUPERPOSITION APPROACH

A common theme in several of the following sections is the superposition approach to thermodynamics, which represents a key component of the computational potential energy landscape framework.¹⁴ Here, we write the total partition function or density of states (DOS) as a sum over contributions from the basins of attraction^{14,32,33} of the local minima, which gives an explicitly ergodic representation of the thermodynamics.^{8,14,34–37} As a first approximation, it is often convenient to use harmonic vibrational densities of states for the local minima. The resulting superposition partition function is nevertheless anharmonic, owing to the distribution of local minima in energy. We can therefore separate contributions to thermodynamic properties in terms of the individual potential wells and the sum over minima. For example, we can identify *well* and *landscape* anharmonicity.^{38,39} For low temperature solid–solid equilibria, the harmonic normal mode approximation can be quantitatively accurate, but methods to sample anharmonicity are generally required if an accurate picture of melting transitions is needed. Reweighting schemes and temperature-dependent frequencies have been considered for this purpose in previous work.^{35,40–43} The most recent basin-sampling scheme⁴⁴ will be outlined in Section *Basin-Sampling*.

The underlying superposition representation for the canonical partition function is

$$\begin{aligned} Z(T) &= \sum_{\alpha=1}^{N^{\text{st}}} \sum_{\zeta=1}^{N_{\alpha}^{\text{PI}}} Z_{\zeta}(T) = \sum_{\alpha=1}^{N^{\text{st}}} N_{\alpha}^{\text{PI}} Z_{\alpha}(T) \\ &= P \sum_{\alpha=1}^{N^{\text{st}}} Z_{\alpha}(T) / o_{\alpha}, \end{aligned} \quad (2)$$

where $Z(T)$, is decomposed in terms of contributions from the catchment basin^{14,32} of each of the N^{st} distinct local minimum structures, with $Z_{\alpha}(T)$ the partition function of structure α at temperature T , which is identical for each of the corresponding N_{α}^{PI} permutation–inversion isomers. Here P is the number of permutation–inversion operations of the Hamiltonian, $2 \prod_{\beta} N_{\beta}!$, and o_{α} is the order of the rigid molecule point group.^{14,45–47}

The great advantage of the superposition approach is that we can exploit efficient methods for locating low-lying local minima based on global optimisation, as described in Section *Basin-Hopping Global Optimisation*. BH approaches that incorporate a local minimisation^{19,48,49} can employ arbitrary moves through configuration space, which can circumvent the barriers that cause ergodicity breaking in standard MC and molecular dynamics procedures. Superposition methods employing both harmonic and anharmonic approximations of the vibrational DOS have found widespread applications in molecular science, and various examples are considered in the following review. First, we outline the BH global optimisation approach that has been used to obtain most of the samples of low-lying potential energy minima. We note that BH is a stochastic procedure, so it is not guaranteed to find all the relevant low-lying minima. However, the sampling of local minima that does not require detailed balance is much faster than for thermodynamic sampling approaches, and is straightforward for the systems considered here.

BASIN-HOPPING GLOBAL OPTIMISATION

The thermodynamic sampling schemes described in the following sections are of particular interest for systems involving broken ergodicity at low temperature, where the potential energy landscape supports alternative low-lying morphologies separated by relatively high barriers. Here, a ‘high’ barrier corresponds to a value that is large compared to $k_{\text{B}}T$ at the temperature, T , where the structures would have equal occupation probabilities, and k_{B} is Boltzmann’s constant. BH global optimisation^{19,48,49} has been successfully employed to survey low-lying minima in a wide range

of systems of this type, including atomic and molecular clusters, biomolecules, and soft and condensed matter.^{14,49,50} These applications will not be reviewed here; since the BH procedure is now well established an overview should be sufficient in the present context.

The key point of the BH framework is to couple local energy minimisation to some sort of step-taking procedure in configuration space. Random perturbations of the atomic coordinates sometimes work quite well, but more efficient schemes can often be devised, which may exploit characteristics of the system under consideration. For example, internal coordinate moves, such as Kirkwood sampling, that respect connectivity are likely to work better for biomolecules with a well-defined covalently bonded framework. Numerous BH variants with alternative step-taking schemes have now been described, and parallel approaches using replicas at different effective temperatures have also been used.⁵¹ In each case, the potential energy for any configuration, \mathbf{X} , becomes the potential energy of the local minimum that the chosen minimisation procedure converges to from \mathbf{X} :

$$\tilde{V}(\mathbf{X}) = \min \{V(\mathbf{X})\}, \quad (3)$$

where \mathbf{X} is a $3N$ -dimensional vector for a system of N atoms. The resulting minimum replaces the previous structure in the chain if it satisfies a chosen acceptance condition, and again several alternatives have been considered. A simple Metropolis scheme often works well, where the new minimum with potential energy E_{new} is accepted if it lies below the potential energy of the starting point, E_{old} . If $E_{\text{new}} > E_{\text{old}}$ it is accepted if $\exp[(E_{\text{old}} - E_{\text{new}})/kT]$ is greater than a random number drawn from the interval $[0,1]$. Schemes based on thresholding,⁵² downhill-only moves,⁵³ and non-Boltzmann weights⁵⁴ based on Tsallis statistics^{55,56} have also been described.

Much larger steps in configuration space can be taken than for conventional molecular dynamics or MC methods, as the energy becomes the value after minimisation, and there is no requirement to satisfy detailed balance. Downhill barriers are therefore removed, and atoms can pass through each other. In fact, there is a more subtle effect, which facilitates sampling when broken ergodicity is prevalent. The occupation probabilities of competing regions of configuration space for the transformed landscape $\tilde{V}(\mathbf{X})$ have been found to overlap significantly over a wider temperature range, where moves still have a good chance of being accepted.^{20,57}

REPLICA-EXCHANGE APPROACHES

The local minima provided by BH can be very helpful to understand systems exhibiting broken ergodicity. However, if thermodynamic properties are required, then configurations within the vicinity of each minimum become important, and increasingly so with higher temperature. At temperatures where well anharmonicity becomes noticeable, alternative techniques are required for sampling the thermodynamic state of interest. In contrast to BH, thermodynamic sampling techniques generally require detailed balance to be satisfied in order to ensure the correct canonical distribution is preserved. While many long-standing strategies exist for sampling stationary distributions, the most notable being MC and thermostatted molecular dynamics algorithms, the restriction of detailed balance often renders such techniques inefficient or impractical. For instance, an MC simulation with random perturbation moves often requires a small displacement step to achieve a useful acceptance probability. As a result, the system only explores its local configuration space on shorter timescales, with large-scale rearrangements, which dominate simulation convergence, occurring on much larger timescales. This quite generic problem necessitates advanced sampling strategies.

REX^{58,59} has emerged as a promising technique to sample such systems exhibiting complex energy landscapes. REX refers to a family of methods in which M -independent copies or 'replicas' of a system are simulated in parallel, typically at different temperatures, with occasional moves that attempt to swap configurations between neighboring replicas. While large energy barriers may prevent the system from exploring its configuration space at low temperatures, enhanced thermal fluctuations provide more rapid barrier crossings at high temperatures. By means of the swap attempts, the low-temperature replicas are exposed to the wider reaches of configuration space explored by the high-temperature replicas, providing a means to sample on both sides of the barrier without directly overcoming it. In this section, we will focus primarily on parallel tempering (PT) in which the replicas are defined by a progression of temperatures, but methods with differing Hamiltonians exist as well.⁶⁰

The REX moves must satisfy detailed balance, which is enforced through an acceptance probability for the exchange attempt. For two replicas, A and B , having the same potential energy function V but at different temperatures T_A and T_B , the acceptance probability takes the form

$$P_{\text{acc}}(\mathbf{X}, \mathbf{Y}) = \min \{1, e^{\Delta\beta\Delta V}\}. \quad (4)$$

Here $\Delta\beta = 1/k_B T_B - 1/k_B T_A$, and $\Delta V = V(\mathbf{Y}) - V(\mathbf{X})$ is the difference in potential energy of the swap configurations \mathbf{X} and \mathbf{Y} in question. Equation (4) guarantees that each replica samples its correct equilibrium state, despite the fact that configurations are being swapped between different thermodynamic ensembles.

Just as detailed balance restricts the perturbation step size in an MC simulation, in REX it limits the temperate spacings between neighboring replicas. To ensure a decent average acceptance rate $\langle P_{\text{acc}} \rangle$, the temperatures T_A and T_B must be spaced closely, yet still far enough apart not to waste computational resources on an unnecessary number of replicas. This conflict has given rise to numerous strategies to optimize various PT parameters, in particular the temperature spacings and total number of replicas.

Understanding a system in terms of the underlying energy landscape can be particularly useful in this respect, and recent developments have been made that enhance PT simulations by utilizing information contained in the energy landscape. These developments, described below, exploit the harmonic superposition approximation (HSA), a harmonic expansion of the landscape about the known configurational minima. Through this approximation, an approximate energy landscape can be constructed solely from a set of minima, and since the form is particularly convenient, namely a series of multidimensional harmonic wells, thermodynamic predictions can be easily made. In this way, thermodynamics from the HSA can be used to approximate the behavior of the true underlying system.

The HSA can be exploited to determine optimal PT temperature spacings, a crucial component in the performance of PT simulations. In general, a uniform acceptance profile is desired across all neighboring replicas,^{61–66} which guarantees that replica round-trip times are not plagued by bottlenecks as trajectories transition between temperatures. Ballard and Wales⁶⁷ have recently demonstrated how knowledge of configurational minima can be utilized to optimize PT temperatures in this respect. By approximating the canonical distribution of each replica via the HSA, they obtained analytic expressions for the average acceptance rate $\langle P_{\text{acc}} \rangle$ in terms of features of the configurational minima. From these expressions, a set of optimal temperatures can be uniquely determined by matching to a target acceptance rate, yielding a progression of replicas that has a uniform acceptance within the approximation of the HSA. Simulations of systems undergoing phase transformation have revealed that this strategy can yield uniform




FIGURE 1 | Acceptance profile for REX simulation of LJ_{31} : The average REX acceptance probability for pair (T_i, T_{i+1}) is plotted versus temperature index i , for temperature spacings chosen using standard geometric progression (empty black circles), and chosen by HSA optimization (filled blue circles). The dip in the geometric case coincides with a heat-capacity peak. At higher temperatures, both profiles deviate from uniform behavior, as the HSA becomes less accurate. Results replotted from Ref 67.

acceptance rates and efficiency enhancements over a standard progression. Figure 1 displays the uniform acceptance profile achieved by the method on PT simulations of LJ_{31} .

In addition to optimizing temperature spacings, the HSA can also be used as a replica in itself. Mandelshtam and coworkers^{7,8} have devised a REX strategy whereby an auxiliary ‘reservoir’ replica is coupled to an M -replica PT simulation. The reservoir replica samples the HSA at a user-specified temperature, with reservoir-temperature swap moves occurring with a particular temperature replica. Because HSA samples can be drawn analytically, reservoir sampling is essentially barrierless and enables rapid exploration of the system’s relevant configuration space (defined by the known minima). In this way, the reservoir reduces the need for high temperature replicas to overcome barriers. Similar strategies for MC sampling employ analogous ideas by generating trial moves near known configurational minima.^{1,5}

KIRKWOOD SAMPLING

Kirkwood sampling⁶⁸ is a method for random (or non-Markovian) sampling of the conformational space of molecules. The principal challenge in random conformational sampling is to avoid steric clashes among the atoms. Kirkwood sampling addresses this problem by incorporating correlations among internal coordinates, as captured by the joint probability

distribution function (pdf) between various sets of internal coordinates. Kirkwood sampling is based on the generalization of the Kirkwood superposition approximation (KSA), originally developed in the radial distribution function theory of liquids.^{69,70} Analogous approximations were later developed^{71,72} in the mutual information theory of correlations and provided expressions for a multidimensional pdf in terms of its marginal pdfs corresponding to neglect of certain correlations. The current application to conformational sampling was motivated by conformational entropy calculation of small molecules using the mutual information expansion of entropy.⁷³ The simplest approximation in this family, namely the KSA, p_3^{KSA} , expresses a three-dimensional pdf, p_3 , in terms of its one- and two-dimensional marginal pdfs

$$p_3(x_1, x_2, x_3) \approx p_3^{\text{KSA}}(x_1, x_2, x_3) = \frac{p_2(x_1, x_2) p_2(x_1, x_3) p_2(x_2, x_3)}{p_1(x_1) p_1(x_2) p_1(x_3)}, \quad (5)$$

where $p_1(\cdot)$ and $p_2(\cdot, \cdot)$ are the marginal pdfs, and the subscript indicates the dimensionality of the pdf. Equation (5) can be obtained from the mutual information expansion of Shannon entropy of p_3 by dropping the threefold mutual information.^{68,73} The KSA can be generalized to express an N -dimensional pdf in terms of its marginals of highest order l .⁷³ For instance, the doublet level ($l=2$) superposition approximation is given by the ratio of the product of 2-D and 1-D marginal pdfs

$$p_N \approx p_N^{(2)}(x_1, \dots, x_N) = \frac{\prod_{i < j} p_2(x_i, x_j)}{\prod_i p_1(x_i)}, \quad (6)$$

where the superscript '(2)' denotes doublet level. Equation 6 accounts for pairwise correlations among all variables but ignores the higher-order correlations. Approximations that account for selected correlations of different orders can also be derived.⁷⁴ Equation 6 is fully coupled and therefore exact sampling would be computationally prohibitive for high-dimensional systems. The Kirkwood approach addresses this issue by sequential sampling of the variables such that each one is selected from a one-dimensional conditional distribution conditioned on the previously sampled variables. The conditional distributions are obtained using the appropriate superposition approximation. For example, given a sequence of variables, the conditional pdf for the k -th variable given values of previous

$k-1$ variables, using the doublet level approximation (Eq. 6), is

$$p_1^{(2)}(x_k | x_1, \dots, x_{k-1}) = \frac{p_k^{(2)}(x_1, \dots, x_{k-1}, x_k)}{p_{k-1}^{(2)}(x_1, \dots, x_{k-1})} = \frac{\prod_{1 \leq j \leq k-1} p_2(x_j, x_k)}{n_k p_1(x_k)^{k-2}}, \quad (7)$$

where the normalization n_k can be computed numerically. The doublet-level Kirkwood scheme effectively samples from the N -dimensional pdf

$$\tilde{p}_N^{(2)}(\vec{x}) = p_2(x_1, x_2) \prod_{3 \leq k \leq N} p_1^{(2)}(x_k | x_1, x_2, \dots, x_{k-1}), \quad (8)$$

and not from Eq. 6. Thus, at the doublet level, the sequence of approximations for the true distribution, p^N , is

$$p_N \rightarrow p_N^{(2)} \rightarrow \tilde{p}_N^{(2)}. \quad (9)$$

Note that both the doublet level Kirkwood approximation, $p_N^{(2)}$, and the doublet-level sampling distribution, $\tilde{p}_N^{(2)}$, involve a product of all singlet and doublet pdfs. Due to this product form, a sampled conformation is guaranteed to fall in the nonzero probability cells of all input pdfs. This construction ensures that all correlations are satisfied simultaneously. The computational complexity of a Kirkwood sampling algorithm is proportional to the number of pdfs used. As the number of doublet-level marginals are $N(N+1)/2$, the doublet-level sampling algorithm has the complexity of $O(N^2)$.

Doublet- and triplet-level Kirkwood sampling has been applied to small drug-like molecules⁶⁸ and small peptides⁷⁵ with up to 52 atoms. In these studies, bond-angle-torsion⁷⁶ internal coordinate system was used and the input pdfs were populated using high-temperature MD simulation data and were effectively marginal pdfs of the Boltzmann distribution at the simulation temperature. Figure 2 shows that the majority of the conformations generated by doublet level sampling for alanine tetrapeptide ($N=150$) had energies similar to those of conformations sampled by the original MD simulation. In other words, the doublet-level Kirkwood sampling distribution is a good approximation of the original Boltzmann distribution. These studies suggest that accounting for just the low-order correlations is sufficient for avoiding steric clashes at the local level, e.g., between adjacent residues of a peptide.

We note that due to high dimensionality and neglect of the higher-order correlations the

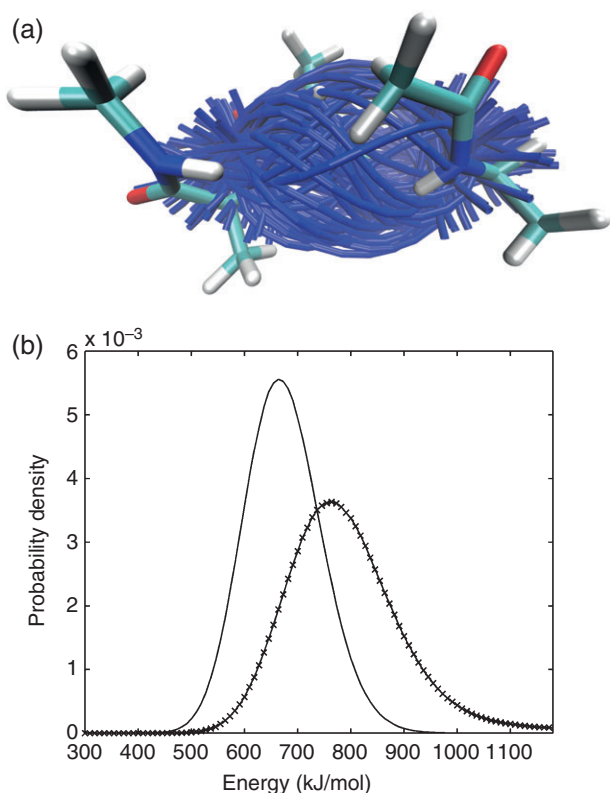


FIGURE 2 | Doublet-level Kirkwood sampling results for 52-atom tetra alanine peptide. The input singlet and doublet pdfs were populated using 5 million conformations from a 500-ns vacuum MD simulation performed at 1000 K. (a) The backbone (blue tube) of 100 Kirkwood sampled conformations aligned on the backbone atoms; one conformation is also shown in licorice. Potential energy for a million Kirkwood samples was computed using the MD energy function. (b) The energy distribution of the Kirkwood samples (crosses) overlaid on the Boltzmann energy distribution obtained from the MD simulation (unmarked). The two distributions have substantial overlap indicating overlap of the Kirkwood sampled and MD-sampled conformational space.

conformational space accessible to Kirkwood sampling will be much larger than that sampled in the original MD simulation. In Figure 2, this observation is reflected in the shift of the Kirkwood energy distribution to higher energies since more conformations are available at higher energies. Unlike MD simulations, Kirkwood sampling generates uncorrelated samples. Furthermore, it is a geometrical sampling method, independent of the potential energy function, and therefore provides barrierless global sampling for any potential energy surface.

In the context of biomolecular simulations, performance of other methods described in this contribution can be enhanced by employing Kirkwood sampling for step taking. For enhanced sampling of local minima, Kirkwood samples could be used

for seeding independent BH simulations. Kirkwood samples can also be used for generating the initial replicas for NS simulations. As the normalized probability of generating a Kirkwood sample is available, a key advantage over knowledge-based conformational samplers,⁷⁷ Kirkwood sampling can also be combined with thermodynamic sampling algorithms that require detailed balance to be satisfied. We next discuss two approaches (Somani et al., unpublished data) for obtaining Boltzmann distributed conformations from Kirkwood sampling.

The first approach utilizes the biased MC¹⁷ framework with the Kirkwood sampling distribution as the biasing distribution. Here Kirkwood samples are used as trial moves and the Metropolis acceptance function is modified to reweight according to the Boltzmann distribution. For example, if doublet-level Kirkwood sampling is used (Eq. 8), the Metropolis acceptance function is given by

$$P_{\text{acc}}(\mathbf{X}_n, \mathbf{X}_o) = \min \left(1, \frac{e^{-\beta U(\mathbf{X}_n; \beta)} / \tilde{p}_N^{(2)}(\mathbf{X}_n)}{e^{-\beta U(\mathbf{X}_o; \beta)} / \tilde{p}_N^{(2)}(\mathbf{X}_o)} \right), \quad (10)$$

where \mathbf{X}_n is the new trial conformation and \mathbf{X}_o is the old conformation. The energy $U(\mathbf{X}_n; \beta)$ includes a contribution due to the Jacobian, $J(\mathbf{X})$, of transformation from internal to Cartesian coordinate system

$$U(\mathbf{X}; \beta) \equiv V(\mathbf{X}) - \frac{1}{\beta} \ln J(\mathbf{X}), \quad (11)$$

where V is the potential energy, typically specified by a molecular mechanics forcefield. In contrast to standard perturbation move MC algorithms that are sequential, the moves in biased MC are independent of the current conformation. Consequently, no equilibration is required and the algorithm can be trivially parallelized in a distributed computing environment. Furthermore, as Kirkwood sampling is a geometrical sampling approach, the same set of Kirkwood samples may be reweighted to generate Boltzmann distributions for different potential energy functions and temperatures.

We have applied (Somani et al., unpublished data) doublet Kirkwood-biased MC simulation to a model system with nine atoms and a bonded chain topology. The input pdfs were generated using data from a 500-K MD simulation and biased MC simulations were performed at successively lower temperatures of 500 K, 400 K, 300 K, and 200 K. Biased MC simulations were able to generate Boltzmann distributions at the MD temperature as well as at lower temperatures, though the acceptance ratios fell with reduced temperatures (Figure 3), consistent with the

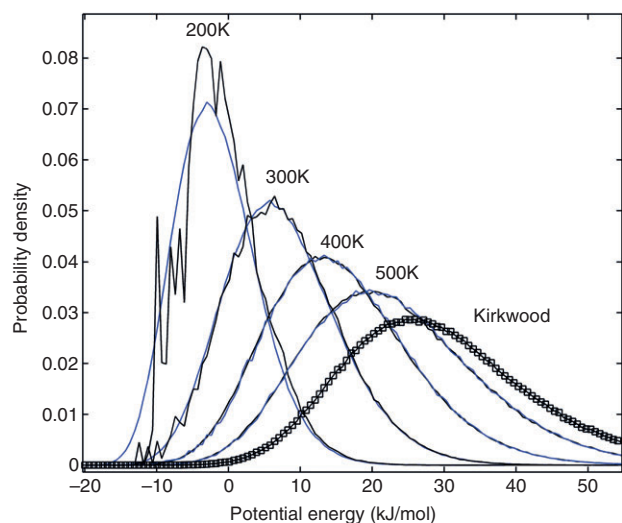


FIGURE 3 | Biased MC simulations for a nine-atom chain molecule using doublet-level Kirkwood sampling distribution as the biasing distribution. Distribution of energy from one million step biased MC simulations (black unmarked line) performed at $T = 500, 400, 300,$ and 200 K, is overlaid on the corresponding reference Boltzmann distributions (in blue unmarked lines). The distribution of energy of the Kirkwood samples is marked by boxes. The acceptance ratio for the different simulations was 0.29, 0.25, 0.07, and 0.009 for $T = 500, 400, 300,$ and 200 K, respectively. These results show that this approach can generate Boltzmann distributions even at temperatures lower than that of the original MD simulation used to populate the input pdfs.

lower overlap. One can imagine an iterative scheme where the initial set of pdfs is constructed in a manner that provides coverage of a wide conformational space. For instance, the pdfs could be populated using high temperature MD, or using a database of conformations from structural databases^{78,79} or a pdf library for molecular fragments. Given this initial set of pdfs and a potential energy function, one could potentially perform successive stages of biased MC simulations and repopulation of the pdfs to reach arbitrarily low temperatures. Note that, for a given potential energy function and temperature, the acceptance ratio of the Kirkwood biased MC simulation is completely determined by the input pdfs and provides a direct measure of the overlap between the Kirkwood sampling distribution and the target Boltzmann distribution. A given acceptance ratio would impose a lower limit on the temperatures for which biased MC simulations can be run.

Due to the product form of the Kirkwood sampling distribution, if there are zero probability cells in the input pdfs then certain regions of the configurational space will not be accessible. As a result Kirkwood sampling will not be ergodic, although the eliminated regions are likely to correspond to conformations with atom clashes. For thermodynamic

sampling, non-ergodicity can be compensated by combining Kirkwood sampling with Markovian samplers, such as MD, or perturbation move MC which explore the conformational space in the vicinity of the current conformation. This exploration is accomplished using the Kirkwood sampler as a reservoir in a reservoir REX simulation. The acceptance function for exchanging between Kirkwood reservoir and a temperature replica is given by

$$P_{\text{acc}}(\mathbf{X}_\beta, \mathbf{X}_r) = \min \left(1, \frac{\tilde{p}_N^{(2)}(\mathbf{X}_\beta)}{\tilde{p}_N^{(2)}(\mathbf{X}_r)} e^{\beta(U(\mathbf{X}_\beta; \beta) - U(\mathbf{X}_r; \beta))} \right), \quad (12)$$

where \mathbf{X}_β is the conformation from a temperature replica and \mathbf{X}_r is reservoir conformation, here drawn from a Kirkwood sampling distribution. The doublet Kirkwood distribution is used in Eq. (12) for illustration. Note that the acceptance function involves the reservoir probability of the conformation from the temperature replica as well as the potential energy of the reservoir conformation. Tests on the nine-atom model system show (Figure 4) improved convergence when a Kirkwood reservoir is employed.

As in the case of biased MC, the acceptance ratio for exchanges with the reservoir is determined by the overlap of the Kirkwood distribution and the Boltzmann distribution for the coupled temperature replica. Note that in the absence of a reservoir the highest temperature in an REX simulation needs to be high enough to overcome the energy barriers and avoid trapping. By coupling to a Kirkwood reservoir, the highest temperature would be dictated by the desired exchange acceptance ratio. Depending on the input pdfs used, Kirkwood dictated highest replica temperature may be lower than the highest temperature dictated by the barriers on the energy landscape. Indeed, if the reservoir has good overlap with the Boltzmann distribution corresponding to the temperature of interest, then just a single replica would suffice. In this case, the temperature replica essentially performs local sampling, while the reservoir facilitates global sampling of the conformational space. One can also imagine a REX simulation where the highest temperature replica is coupled to a Kirkwood reservoir to facilitate global sampling and the lowest temperature is coupled to an HSA reservoir constructed from the low-energy minima to enable rapid equilibration over the low-energy regions of the energy landscape.

SUPERPOSITION-ENHANCED NESTED SAMPLING

NS is an importance sampling method that was recently introduced in the Bayesian statistics




FIGURE 4 | Potential energy distributions (lines marked by circles) from the 5×10^6 steps temperature replica exchange simulations (a) without and (b) with a doublet-level Kirkwood reservoir. The replica temperatures are 20 (black), 30 (red), 50 (blue), and 100 K (green). The distributions from the Kirkwood reservoir simulation are much closer to the reference distributions (unmarked lines), indicating improved convergence upon coupling to a reservoir.

community⁸⁰ and has since been applied in a variety of fields, including astrophysics,^{81,82} systems biology,⁸³ and statistical physics.^{84–86} NS was originally described in the language of Bayesian statistics, but can equally be applied to thermodynamic sampling. The Bayesian prior is the uniform distribution over all of phase space (every point in phase space equally likely). The integral over phase space is performed, weighted by the likelihood, which is simply the Boltzmann weight $\exp(-V(\mathbf{X})/T)$. In the execution of NS, the value of the likelihood is unimportant; only the ordering matters, which means that the energy can be used as a proxy for the Boltzmann weight, making NS independent of temperature. The primary output of NS is the DOS, from which it is possible to compute thermodynamic quantities at any temperature.

Nested sampling proceeds by iteratively constructing a list of energy levels $\{E_1^{\max}, \dots, E_N^{\max}\}$ which are sorted, so that $E_i^{\max} > E_{i+1}^{\max}$. These energy levels have the special property that the volume of phase space with energy below E_i^{\max} satisfies, on average

$$\Omega(E < E_i^{\max}) = \alpha \Omega(E < E_{i-1}^{\max}), \quad (13)$$

where $0 < \alpha < 1$ is determined by a parameter of the algorithm (Eq. 15). From this result, the density of states, g_i , the normalized phase space volume with energy between E_i^{\max} and E_{i-1}^{\max} is constructed as

$$g_i = \frac{\Omega(E < E_{i-1}^{\max}) - \Omega(E < E_i^{\max})}{\Omega_0} = \alpha^{i-1} - \alpha^i, \quad (14)$$

where Ω_0 is the total phase space volume.

The NS algorithm begins by generating K configurations of the system (replicas) distributed randomly and uniformly within phase space Ω_0 . For systems with infinite phase space volume the system can be placed in a box. Alternatively, one can limit phase space by placing an upper bound on the energy. Either way, the actual value of Ω_0 is only a normalization constant and does not come into the definition of the DOS. The energies of the replicas are computed and stored in a list $\{E_1, \dots, E_K\}$. The key insight of NS, is that the energy of the replica with the maximum energy (E_1^{\max}) satisfies equation 13 as $\Omega(E < E_1^{\max}) \approx \alpha \Omega_0$ with the compression factor

$$\alpha = \frac{K}{K+1}. \quad (15)$$

This result can be understood by laying out every point in phase space on a one-dimensional line and sorting them by energy. The ‘volume’ of the line is just the length, and goes from 0 (leftmost) to Ω_0 (rightmost). The K replicas are distributed randomly and uniformly on that line. The replica with the largest energy will be the rightmost replica. If K points are placed on a unit line, the position of the point with the largest value will be distributed according to

$$p_K(x) \sim x^{K-1}, \quad (16)$$

with expectation value $K/(K+1)$. Thus, the mean position of the right-most replica is $\Omega_0 K/(K+1)$, that also corresponds to $\Omega(E < E_1^{\max})$.

We now remove the replica with the maximum energy and replace it with a new replica sampled uniformly from with the volume $\Omega(E < E_1^{\max})$. This is typically done by initializing the new replica at the position of one of the others and walking it via an MC Markov chain for sufficient steps that it loses its memory of the starting location. We now again have K replicas sampled uniformly from the volume $\Omega(E < E_1^{\max})$. Thus, using the same arguments as before, the energy of the maximum energy replica is saved as E_2^{\max} , and satisfies Eq. 13. This procedure is iterated until some stopping condition is met (Figure 5).

A major benefit of NS is that, although phase space is divided into energy bins, as in Eq. 14, the energy bins are determined adaptively via the constant compression factor $\alpha = K/(K+1)$. This stepping procedure automatically creates higher resolution for low-energy regions of phase space and in regions where phase space volume is changing most rapidly, which usually applies near phase transitions. This feature of adaptively determined bin sizes is in contrast with the energy bins in the Wang-Landau method,³ and the temperature spacing in PT, which must be predetermined, although a method for choosing the temperature spacing was presented in Section *Replica-Exchange Approaches*. In fact, Brewer et al.⁸⁶ introduced diffusive NS, which uses the energy bins and weights g_i produced by NS as input to the Wang-Landau method to refine the estimate for the DOS.

The vast majority of the computation time in NS is spent generating new configurations sampled uniformly from the space $\Omega(E < E_i^{\max})$. As discussed above, the standard way of achieving this goal is using an MC walk long enough so that the replica loses memory of where it started. Uniform sampling is maintained by accepting every step subject only to the criterion that the new energy is less than E_i^{\max} .

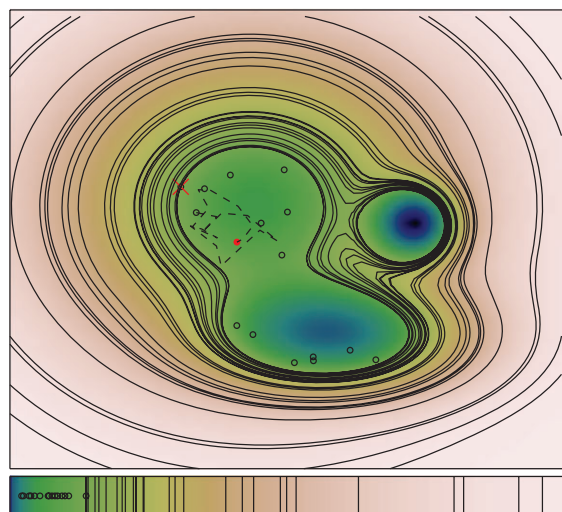


FIGURE 5 | The nested sampling procedure is illustrated for a two-dimensional energy landscape where darker colors indicate lower energies. The image is drawn after 26 iterations of the nested sampling procedure. The circles represent the replica positions ($K = 15$), which are distributed uniformly in the two-dimensional space. The lines are constant energy contours at E_i^{\max} for each iteration i . The cross gives the location of the highest energy replica, which defines E^{\max} for the next iteration. The dotted line is the path of the MC walk of the newly generated replica, which is shown as a solid circle. The lower panel shows the one-dimensional representation of phase space, which is used to derive the compression factor of Eqs 13 and 15. In this configuration, there are no replicas in the basin with the global minimum, which will cause the density of states to be overestimated. This major drawback of the NS algorithm can be overcome by superposition enhanced nested sampling if the basin-hopping procedure correctly identifies the relevant low-lying minima.

This technique can lead to serious sampling problems for multimodal energy landscapes. If two minima have energy less than E_i^{\max} , but the minimum energy path between them lies above E_i^{\max} , then an MC walk started in one basin will never reach the other. This problem is partially alleviated by the fact that there are many independent replicas. As long as there are sufficient replicas in each basin then the statistics should not be affected too much. However if, through statistical fluctuations, one basin is devoid of replicas, it will never become repopulated. This issue is most likely to cause serious problems in systems where the basin with the global minimum is narrow and separated by large energy barriers from the rest of phase space. Examples of such difficult systems are LJ₃₁, LJ₃₈, and LJ₇₅.⁸⁵

New methods for uniform sampling with a likelihood (or energy) constraint are actively being developed.⁸⁶ Here we describe a recent method that uses ideas from potential energy landscape theory to speed up sampling and overcome the problem of being




FIGURE 6 | Heat capacity curves for LJ₃₁. HSA corresponds to the harmonic superposition approximation. All NS and SENS calculations were performed using $K = 20,000$ replicas.

locked out of basins, namely superposition-enhanced NS (SENS).⁸⁷

SENS combines the strength of global optimisation algorithms with that of NS to achieve an augmented exploration of the low energy regions of the energy landscape, while maintaining all the desirable properties of NS. Global optimisation algorithms, such as BH (Section *Basin-Hopping Global Optimisation*), are much more efficient at finding the low energy configurations by virtue of a greater freedom to use energy minimisation algorithms and MC moves that do not respect detailed balance. The collection of minimum energy configurations thus obtained can be used to augment the exploration of the energy landscape performed by NS. SENS contributes knowledge of this collection of low-energy configurations via Hamiltonian REX (HREM).^{88,89} Similar in spirit to PT, HREM moves are proposed between configurations sampled using different Hamiltonians (energy functions), rather than sampling configurations with the same Hamiltonian, but weighted at different temperatures, as in PT.

In this particular instance, similar in spirit to Ref 89, a minimum is selected from the precomputed collection of minima weighted by their HSA configurational volume and a swap is proposed between a configuration X_{har} , sampled uniformly within its

harmonic basin of attraction, and a configuration X_{sys} , sampled by an MC walk using the system Hamiltonian under a certain energy constraint E_{max} . If

$$H_{\text{har}}(X_{\text{sys}}) \leq E_{\text{max}} \quad \text{and} \quad H_{\text{sys}}(X_{\text{har}}) \leq E_{\text{max}}, \quad (17)$$

then the true and the HSA distribution of states overlap and the swap is accepted. HREM moves are typically followed by conventional MC walks to further explore neighboring regions of configurational space. The acceptance of such HREM moves depends heavily on how well the HSA approximates the underlying potential energy surface and so does the efficiency of SENS over that of NS. While the HSA captures landscape anharmonicity, arising from the distribution of local minima, it does not include well anharmonicity, thus becoming an increasingly good approximation at lower energies. For this reason, we expect that only the low-energy minima will contribute significantly to augmenting the sampling.

An example of a very successful application of SENS is found for the thermodynamic analysis of LJ₃₁, where the method has been shown to be at least one order of magnitude more efficient than PT and NS.⁸⁷ In Figure 6, we show results for the heat capacity curve of LJ₃₁; for $K = 20,000$ replicas SENS is well converged and agrees with the HSA at low temperatures, while for NS the low-temperature peak is completely absent.




FIGURE 7 | Comparison of heat-capacity curves for LJ₃₁ obtained by exact SENS using different numbers of replicas. The PT and HSA curves were obtained by parallel tempering and the harmonic superposition approximation, respectively. Figure reproduced from Ref 87.

In Figure 7, we show a comparison of PT, HSA, and exact SENS for a range of replica numbers (see Table 1 for details). SENS can accurately reproduce the low temperature features of the heat capacity for as low as $K = 2500$ replicas, representing an improvement in performance of 20 times over PT.⁸⁷ The benefits are not as great for LJ₇₅ where even reaching the region where the HSA becomes accurate enough for swaps to be accepted is challenging.

An approximate version of the method has been proposed to alleviate this problem: similar in spirit to the basin-sampling approach (Section *Basin Sampling*), approximate-SENS interpolates between the high energy DOS obtained by NS and the low-energy DOS provided by the HSA. Once at low energy, this interpolation can be achieved by starting a fraction of the MC walks from a local minimum configuration sampled from the precomputed collection of minima weighted by their HSA volume. The transition from the conventional sampling at high energy to the enhanced sampling at low energy is controlled by an onset function. The choice of the energy at which the transition should occur is practically the only additional parameter when compared to exact SENS and NS. Approximate-SENS is easy to implement and generally yields equally good or

TABLE 1 | Comparison of Methods Used to Obtain the LJ₃₁ Heat-Capacity Curves Shown in Figures 6 and 7

LJ ₃₁			
Method	K	N	$N_E^{(\text{total})}$
PT			1.9×10^{11}
NS ⁸⁵	280,000		3.4×10^{12}
SENS approx	20,000	10,000	1×10^{11}
SENS exact	20,000	10,000	1×10^{11}
SENS exact	10,000	10,000	5.2×10^{10}
SENS exact	5000	10,000	2.6×10^{10}
SENS exact	2500	10,000	1.3×10^{10}

$N_E^{(\text{total})}$ indicates the total number of energy evaluations (summed over all processors). PT was performed using 24 replicas spread geometrically through the temperature range of 0.0125–0.6. Note that, although not shown, approximate SENS can perform as well as exact SENS when fewer replicas are used, results for LJ₇₅ illustrate clearly the capabilities of the method.⁸⁷ $N_E^{(\text{tot})} = N \times \mathcal{P} \times N_{\text{iter}}$ is the total number of energy evaluations, where N is the length of each Monte Carlo walk, \mathcal{P} is the number of processors used, and N_{iter} is the total number of nested sampling iterations.

better results than exact SENS, although it is formally biased.

The thermodynamic analysis of LJ₇₅ by approximate-SENS converges in about $O(10^{11})$

energy evaluations, unlike PT, which has been shown to never converge on conventional simulation timescales. The SENS methods are therefore another viable solution to tackle systems exhibiting broken ergodicity.

BASIN-SAMPLING

The recently described⁴⁴ basin-sampling procedure is perhaps the most efficient scheme currently available for sampling global thermodynamic properties. This approach combines knowledge of low-lying local minima on the potential energy landscape with PT, to connect densities of states that are relevant in the low and high temperature regimes. A two-dimensional histogram of probabilities is constructed using the instantaneous potential energy and the energy of local minima obtained by regular quenching. An approximate anharmonic functional form is then fitted to the results corresponding to energy bins for the distribution of local minima on the landscape. This construction has the additional benefit that it can be employed

to calculate the potential energy density of both local minimum structures and permutation–inversion isomers, as well as estimates of the total number of structures and isomers on the landscape.

The regular quenching that does not affect the instantaneous coordinates in the Markov chain provides statistics for the number of visits to quench potential energy bin q with potential energy V_q^Q from instantaneous potential energy bin i in replica r , denoted by \mathcal{N}_{igr} . The canonical probability distribution

$$\bar{P}(V_i^I, V_q^Q, T_r) = \mathcal{N}_{igr} / \mathcal{N}_r \propto \bar{\Omega}_c(V_i^I, V_q^Q) e^{-V_i^I / k_B T_r}, \quad (18)$$

is then obtained by minimising

$$\chi_{2D}^2 = \sum_r \sum_i \sum_q \mathcal{N}_{igr} \left[\ln \bar{\Omega}_c(V_i^I, V_q^Q) - \ln \left(\frac{\mathcal{N}_{igr} e^{V_i^I / k_B T_r} Z_c(T_r)}{\mathcal{N}_r} \right) \right]^2, \quad (19)$$

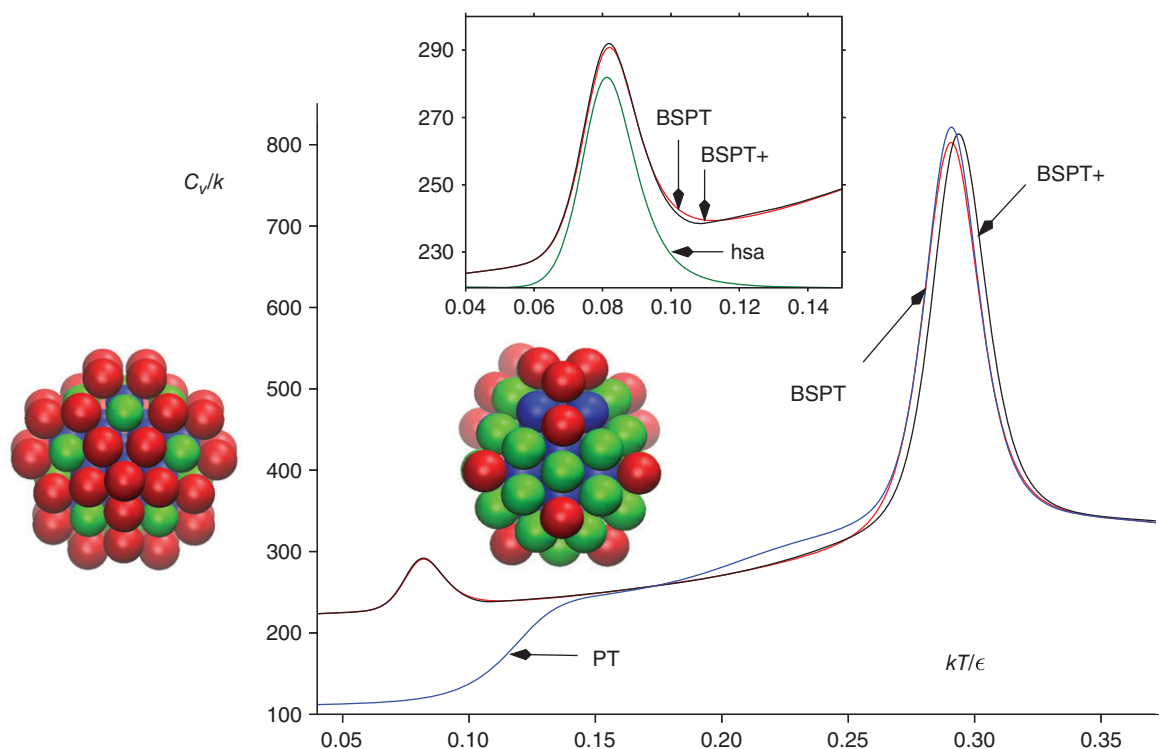


FIGURE 8 | Heat capacity as a function of reduced temperature for LJ₇₅.⁴⁴ The lowest and second-lowest minima based on a Marks decahedron and an incomplete Mackay icosahedron are illustrated on the left and right of the low-temperature heat-capacity peak, respectively. The atoms are coloured according to their contribution to the total energy: the most tightly bound atoms are blue, the least tightly bound are red, with intermediate binding energies in green. The curve marked PT is for the parallel tempering data only from the BSPT run. The inset shows a magnification of the low-temperature peak corresponding to the solid–solid transition and a comparison with the harmonic superposition result. The curves marked BSPT+ result when the original statistics are combined with longer runs aimed at converging the potential energy densities of local minima, and provide a consistency check.




FIGURE 9 | LJ₇₅ probability distribution $P(Q_6, T)$ and the corresponding free energy surface $F(Q_6, T) = -k_B T \ln P(Q_6, T)$ for a database containing 8,391,630 structures.⁴⁴ Four structures are illustrated, namely the global minimum (decahedron, $Q_6 = 0.31$), the lowest minima based upon icosahedral packing with anti-Mackay and Mackay overlayers⁹⁴ ($Q_6 = 0.02$ and 0.15 , respectively), and a minimum associated with the liquid-like phase ($Q_6 = 0.10$).

where the variables are $\bar{\Omega}_c(V_i^I, V_q^O)$ if we fix $Z_c(T_r) = Z_c^*(T_r)$ from an initial 1D fit to the PT results for the distribution $P(V, T) \propto \Omega_c(V) e^{-V/k_B T}$. Here $\Omega_c(V)$ is again the configurational DOS. For each replica r with temperature T_r , we count the number of visits, \mathcal{N}_{ir} , to potential energy bins indexed as V_i^I , providing the estimate for $P(V_i^I, T_r)$ as

$$P(V_i^I, T_r) = \mathcal{N}_{ir} / \mathcal{N}_r, \quad (20)$$

where $\mathcal{N}_r = \sum_i \mathcal{N}_{ir}$ is the total number of MC steps for replica r . We then minimise

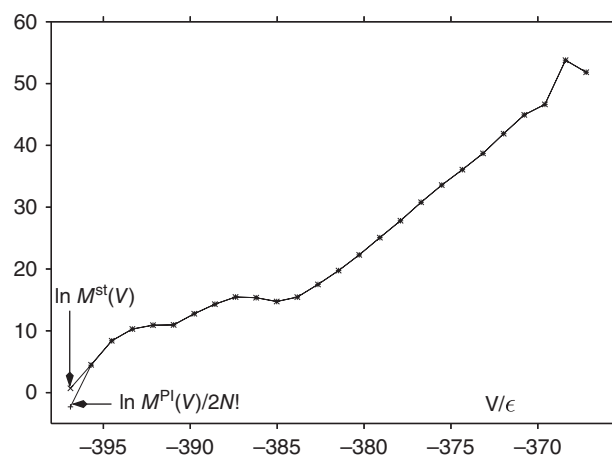


FIGURE 10 | Calculated values for the $\ln M^{\text{st}}(V)$ and $\ln M^{\text{pl}}(V)/2N!$ as a function of potential energy for LJ₇₅.⁴⁴ $M^{\text{st}}(V)$ and $M^{\text{pl}}(V)$ are the potential energy densities of distinct local minima structures and permutation–inversion isomers, respectively.

$$\chi_{1\text{D}}^2 = \sum_r \sum_i \mathcal{N}_{ir} \left[\ln \Omega_c(V_i^I) - \ln \left(\frac{\mathcal{N}_{ir} e^{V_i^I/k_B T_r} Z_c(T_r)}{\mathcal{N}_r} \right) \right]^2, \quad (21)$$

where the variables are $\Omega_c(V_i^I)$ and $Z_c(T_r)$. The optimal values are denoted by $\Omega_c^*(V_i^I)$ and $Z_c^*(T_r)$. In this work, all the fitting was conducted by minimising $\chi_{2\text{D}}^2$ or $\chi_{1\text{D}}^2$ using the modified limited memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm^{90,91} from the GMIN code.⁹²

To simplify the fitting, a model anharmonic DOS was employed for each quench bin, and an efficient representation was obtained using two fitting parameters, A_q and B_q , with

$$\ln \bar{\Omega}_c(V_i^I, V_q^O) = (\kappa' + e^{A_q} V_{iq}) \ln V_{iq} + B_q, \quad (22)$$

where $\kappa' = \kappa/2 - 1$ and $V_{iq} = V_i^I - V_q^O$, the difference between the instantaneous and quench potential energies. Here we have built on previous work⁴⁰ that used the analytic DOS for a Morse potential.⁹³ For $V_{iq} \rightarrow 0$, we recover the usual harmonic result: $\bar{\Omega}_c(V_i^I, V_q^O) \propto V_{iq}^{\kappa/2-1}$. A_q makes no contribution in this limit, and can therefore be identified with the effective well anharmonicity, while B_q incorporates the landscape entropy in terms of the number of minima included in quench bin q . This interpretation facilitates the calculation of distributions for the potential energy density of minimum energy structures.⁴⁴

Optimal values for the two parameters in each bin were again obtained by fitting, and are denoted A_q^*

and B_q^* . Then to exploit additional information corresponding to the low temperature limit a normal mode analysis of the local minima obtained from BH global optimisation was used. B_q^* values were simply replaced up to a specified potential energy threshold using the normal mode data, and the final configurational DOS, $\Omega_c(V_i^l)$, was obtained by combining the results of the one- and two-dimensional fitting procedures as:

$$\Omega_c(V_i^l) = \sum_q \left[\delta_q \Omega_c^*(V_i^l) \frac{\mathcal{N}_{iq}}{\mathcal{N}_i} + (1 - \delta_q) \overline{\Omega}_c^*(V_i^l, V_q^Q) \right], \quad (23)$$

where the mixing parameter $\delta_q = (i - i_{\min}^q) / (i_{\max}^q - i_{\min}^q)$ depends on the optimal fitting range for the q bins.⁴⁴

Extensive tests were conducted for both LJ₃₁ and LJ₇₅. Some results for the larger cluster are illustrated in Figures 8–10. Here Q_6 is a bond-order parameter,^{95,96} which takes larger values of the decahedral global minimum than for minima based upon icosahedral packing. Thirty-two temperature replicas were used, exponentially spaced in the temperature range of 0.15–0.375. The production run of 200×10^6 standard PT steps followed by 60×10^6 BSPT steps (PT with quenches every 30 steps) required 22 h of wall clock time. When the lowest 13 minima are used to replace the fitted B_q^* values the resulting heat capacity curve appears to be well converged (Figure 8, with the characteristic low-temperature solid–solid peak well reproduced. This result can be compared with previous simulations using an auxiliary harmonic superposition reference,⁷ which required 3×10^9 steps. In contrast, even 10^{11} MC steps are not enough to converge the heat capacity using adaptive exchange PT.^{8,89} The corresponding probability distribution $P(Q_6, T)$ and free-energy

surface are shown in Figure 9, where various features corresponding to different families of structures can be associated with peaks that are separated in these two-dimensional projections. The calculated potential energy distributions of distinct structures and permutation-inversion isomers are shown in Figure 10. Here the effect of point group symmetry in reducing the number of permutation–inversion isomers is visible at low energy. There is also an unexpected feature, namely a shallow minimum in the distributions around $V = -385 \epsilon$, which has been associated with regions of the potential energy landscape that are relatively sparsely populated.⁴⁴ The total number of distinct local minimum structures for this cluster is estimated at around 4×10^{25} .⁴⁴ The capability of basin-sampling to yield estimates like this as a by-product of the sampling could be particularly interesting for amorphous systems in future studies.

CONCLUSIONS

The idea of exploiting knowledge of low-lying local minima to enhance sampling^{1,5} recently reached fruition in a variety of new techniques.^{7,38,44,67,87} The foundation for each of these methods is the superposition approach, in which the total partition function is decomposed into contributions from distinct local minima.^{8,14,34–37} The continuing development of these methods provides important cross-validations of results for challenging systems that exhibit broken ergodicity. These results are not only important in themselves, but also provide an essential platform for optimising efficiency, which is likely to come from hybrid schemes, and may be system dependent. Nevertheless, many general conclusions should hold for very diverse problems in chemical physics, and important tests in molecular simulation and soft and condensed matter problems should reveal that further progress is possible.

ACKNOWLEDGMENTS

The authors gratefully acknowledge financial support from the EPSRC and the ERC. S.M. acknowledges financial support from the Gates Cambridge Scholarship. We also thank Gábor Csányi for helpful discussions.

REFERENCES

1. Zhou R, Berne BJ. Smart walking: a new method for Boltzman sampling of protein conformations. *J Chem Phys* 1997, 107:9185–9196.
2. Hansmann UHE, Okamoto Y. *Curr Opin Struct Biol* 1999, 9:177–183.
3. Wang F, Landau DP. *Phys Rev Lett* 2001, 86:2050–2053.
4. Watanabe W, Reinhardt WP. *Phys Rev Lett* 1990, 65:3301–3304.

5. Andricioaei I, Straub JE, Voter AF. Smart darting Monte Carlo. *J Chem Phys* 2001, 114:6994–7000.
6. Kim J, Straub JE, Keyes T. Statistical-temperature Monte Carlo and molecular Dynamics Algorithms. *Phys Rev Lett* 2006, 97:050601.
7. Sharapov VA, Mandelshtam VA. Solid-solid structural transformations in Lennard-Jones clusters: accurate simulations versus the harmonic superposition approximation. *J Phys Chem A* 2007, 111:10284–10291.
8. Sharapov VA, Meluzzi D, Mandelshtam VA. Low-temperature structural transitions: Circumventing the broken-ergodicity problem. *Phys Rev Lett* 2007, 98:105701.
9. Kim J, Keyes T, Straub JE. Communication: Iteration-free, weighted histogram analysis method in terms of intensive variables (4 pages). *J Chem Phys* 2011, 135:061103.
10. Kim J, Straub JE, Keyes T. Replica exchange statistical temperature molecular dynamics algorithm. *J Phys Chem B* 2012, 116:8646–8653.
11. Ortoleva PJ, Keyes T, Tuckerman M. Macromolecular systems understood through multiscale and enhanced sampling techniques. *J Phys Chem B* 2012, 116:8335–8336.
12. Huber GA, Kim S. Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophys J* 1996, 70:97–110.
13. Bhatt D, Zhang BW, Zuckerman DM. Steady-state simulations using weighted ensemble path sampling (12 pages). *J Chem Phys* 2010, 133:014110.
14. Wales DJ. *Energy Landscapes*. Cambridge: Cambridge University Press; 2003.
15. Wales DJ. Energy landscapes: some new horizons. *Curr Opin Struct Biol* 2010, 20:3–10.
16. Wales DJ. Decoding the energy landscape: extracting structure, dynamics and thermodynamics. *Philos Trans R Soc A* 2012, 370:2877–2899.
17. Frenkel D, Smit B. *Understanding Molecular Simulation*. 2nd ed. London: Academic Press; 2002.
18. Jones JE, Ingham AE. On the calculation of certain crystal potential constants, and on the cubic crystal of least potential energy. *Proc R Soc Lond A* 1925, 107:636–653.
19. Wales DJ, Doye JPK. Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms. *J Phys Chem A* 1997, 101:5111–5116.
20. Doye JPK, Wales DJ, Miller MA. Thermodynamics and the global optimization of lennard-jones clusters. *J Chem Phys* 1998, 109:8143–8153.
21. Doye JPK, Miller MA, Wales DJ. The double-funnel energy landscape of the 38-atom Lennard-Jones cluster. *J Chem Phys* 1999, 110:6896–6906.
22. Calvo F, Neirotti JP, Freeman DL, Doll JD. Phase changes in 38 atom Lennard-Jones clusters, II: a parallel tempering study of equilibrium and dynamic properties in the molecular dynamics and microcanonical ensembles. *J Chem Phys* 2000, 112:10350–10357.
23. Neirotti JP, Calvo F, Freeman DL, Doll JD. Phase changes in 38 atom Lennard-Jones clusters, I: a parallel tempering study in the canonical ensemble. *J Chem Phys* 2000, 112:10340–10349.
24. Frantsuzov PA, Mandelshtam VA. Size-temperature phase diagram for small Lennard-Jones clusters (4 pages). *Phys Rev E* 2005, 72:037102.
25. Predescu C, Frantsuzov PA, Mandelshtam VA. Thermodynamics and equilibrium structure of Ne₃₈ cluster: quantum mechanics versus classical (12 pages). *J Chem Phys* 2005, 122:154305.
26. Liu H, Jordan KD. Finite temperature properties of (CO₂)_n clusters. *J Phys Chem A* 2003, 107:5703–5709.
27. Wille LT. Lennard-Jones clusters and the multiple-minima problem. In: Stauffer D, ed. *Annual Reviews of Computational Physics VII*. Singapore: World Scientific; 2000.
28. Wales DJ. Discrete path sampling. *Mol Phys* 2002, 100:3285–3305.
29. Dellago C, Bolhuis PG, Chandler D. Efficient transition path sampling: application to Lennard-Jones cluster rearrangements. *J Chem Phys* 1998, 108:9236–9245.
30. Passerone D, Parrinello M. Action-derived molecular dynamics in the study of rare events. *Phys Rev Lett* 2001, 87:108302.
31. Ren WEW, Vanden-Eijnden E. *Phys. Rev. B* 2002, 66:052301.
32. Mezey PG. Catchment region partitioning of energy hypersurfaces, I. *Theor Chim Acta* 1981, 58:309–330.
33. Mezey PG. *Potential Energy Hypersurfaces*. Amsterdam: Elsevier; 1987.
34. Stillinger FH, Weber TA. Packing structures and transitions in liquids and solids. *Science* 1984, 225:983–989.
35. Wales DJ. Coexistence in small inert gas clusters. *Mol Phys* 1993, 78:151–171.
36. Stillinger FH. A topographic view of supercooled liquids and glass formation. *Science* 1995, 267:1935.
37. Strodel B, Wales DJ. Free energy surfaces from an extended harmonic superposition approach and kinetics for alanine dipeptide. *Chem Phys Lett* 2008, 466:105–115.
38. Bogdan TV, Wales DJ, Calvo F. Equilibrium thermodynamics from basin-sampling (13 pages). *J Chem Phys* 2006, 124:044102.
39. Sciortino F, Kob W, Tartaglia P. Thermodynamics of supercooled liquids in the inherent-structure formalism: a case study. *J Phys Condens Mat* 2000, 12:6525–6534.
40. Doye JPK, Wales DJ. Calculation of thermodynamic properties of small Lennard-Jones clusters incorporating an harmonicity. *J Chem Phys* 1995, 102:9659–9672.

41. Doye JPK, Wales DJ. An order parameter approach to coexistence in atomic clusters. *J Chem Phys* 1995, 102:9673–9688.
42. Calvo F, Doye JPK, Wales DJ. Characterization of anharmonicities on complex potential energy surfaces: perturbation theory and simulation. *J Chem Phys* 2001, 115:9627–9636.
43. Georgescu I, Mandelstam VA. Self-consistent phonons revisited, I: the role of thermal versus quantum fluctuations on structural transitions in large Lennard-Jones clusters. *J Chem Phys* 2012, 137:144106.
44. Wales DJ. Surveying a complex potential energy landscape: overcoming broken ergodicity using basin-sampling. *Chem Phys Lett* 2013, 584:1–9.
45. Amar FG, Berry RS. The onset of non-rigid dynamics and the melting transition in Ar₇. *J Chem Phys* 1986, 85:5943–5954.
46. Gilson MK, Irikura KK. Symmetry numbers for rigid, flexible, and fluxional molecules: theory and applications. *J Phys Chem B* 2010, 114:16304–16317.
47. Calvo F, Doye JPK, Wales DJ. Energy landscapes of colloidal clusters: thermodynamics and rearrangement mechanisms. *Nanoscale* 2012, 4:1085–1100.
48. Li Z, Scheraga HA. Monte Carlo-minimization approach to the multiple-minima problem in protein folding. *Proc Natl Acad Sci USA* 1987, 84:6611–6615.
49. Wales DJ, Scheraga HA. Global optimization of clusters, crystals and biomolecules. *Science* 1999, 285:1368–1372.
50. Klenin K, Strodel B, Wales DJ, Wenzel W. Modelling proteins: conformational sampling and reconstruction of folding kinetics. *Biochim Biophys Acta* 2011, 1814:977–1000.
51. Strodel B, Lee JW, Whittleston CS, Wales DJ. Transmembrane structures for Alzheimer's A β (1-42) oligomers. *J Am Chem Soc* 2010, 132:13300–13312.
52. Hoffmann KH, Franz A, Salamon P. Structure of best possible strategies for finding ground states. *Phys Rev E* 2002, 66:046706.
53. Leary RH, Doye JPK. New tetrahedral global minimum for the 98-atom Lennard-Jones cluster. *Phys Rev E* 1999, 60:R6320–R6322.
54. Shang C, Wales DJ. Communication: optimal parameters for basin-hopping global optimization based on Tsallis statistics. *J Chem Phys* 2014, 141:071101.
55. Tsallis C. *J Stat Phys* 1988, 52:479–487.
56. Tsallis C. Nonextensive statistics: theoretical, experimental and computational evidences and connections. *Braz J Phys* 1999, 29:1–35.
57. Doye JPK, Wales DJ. Thermodynamics of global optimization. *Phys Rev Lett* 1998, 80:1357–1360.
58. Swendsen RH, Wang J-S. Replica Monte Carlo simulation of spin-glasses. *Phys Rev Lett* 1986, 57:2607–2609.
59. Earl D, Deem MW. Parallel tempering: theory, applications, and new perspectives. *Phys Chem Chem Phys* 2005, 7:3910–3916.
60. Sugita Y, Okamoto Y. Replica-exchange multicanonical algorithm and multicanonical replica-exchange method for simulating systems with rough energy landscape. *Chem Phys Lett* 2000, 329:261–270.
61. Sugita Y, Okamoto Y. Replica-exchange molecular dynamics method for protein folding. *Chem Phys Lett* 1999, 314:141–151.
62. Kofke DA. On the acceptance probability of replica-exchange Monte Carlo trials. *J Chem Phys* 2002, 117:6911–6914.
63. Kofke DA. Erratum: “on the acceptance probability of replica-exchange Monte Carlo trials. *J Chem Phys* 2004, 120:10852.
64. Nadler W, Hansmann UHE. Generalized ensemble and tempering simulations: a unified view. *Phys Rev E* 2007, 75:026109.
65. Nadler W, Hansmann UHE. Optimized explicit-solvent replica exchange molecular dynamics from scratch. *J Phys Chem B* 2008, 112:10386–10387.
66. Denschlag R, Lingenheil M, Tavan P. Optimal temperature ladders in replica exchange simulations. *Chem Phys Lett* 2009, 473:193–195.
67. Ballard AJ, Wales DJ. Superposition-enhanced estimation of optimal temperature spacings for parallel tempering simulations. *J Chem Theory Comput* 2014, 10:5599–5605.
68. Somani S, Killian BJ, Gilson MK. Sampling conformations in high dimensions using low-dimensional distribution functions. *J Chem Phys* 2009, 130:34102.
69. Kirkwood JG, Boggs EM. The radial distribution function in liquids. *J Chem Phys* 1942, 10:394–402.
70. Hansen J, McDonald I. *Theory of Simple Liquids*. 3rd ed. Amsterdam: Academic Press; 1976.
71. McGill W. Multivariate information transmission. *IEEE Trans Inf Theory* 1954, 4:93–111.
72. Fano RM. *Transmission of Information: A Statistical Theory of Communication*. New York: The MIT Press; 1961.
73. Killian BJ, Gilson MK, Kravitz JY. Extraction of configurational entropy from molecular simulations via an expansion approximation. *J Chem Phys* 2007, 127:24107.
74. Somani S. Conformational sampling and calculation of molecular free energy using superposition approximations. PhD Thesis, University of Maryland, College Park, MD, 2011.
75. Somani S, Gilson MK. Accelerated convergence of molecular free energy via superposition approximation-based reference states. *J Chem Phys* 2011, 134:34107.

76. Chang C-E, Potter MJ, Gilson MK. Calculation of molecular configuration integrals. *J Phys Chem B* Jan. 2003, 107:1048–1055.
77. Chen I-J, Foloppe N. Is conformational sampling of drug-like molecules a solved problem? *Drug Dev Res* 2011, 72:85–94.
78. Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE. The protein data bank. Available at: <http://www.rcsb.org/pdb>. (Accessed January 6, 2014).
79. Allen FH. The Cambridge Structural Database: a quarter of a million crystal structures and rising. *Acta Cryst* 2002, 58:380–388.
80. Skilling J. Nested sampling for general Bayesian computation. *Bayesian Anal* 2006, 1:833–859.
81. Vigeland SJ, Vallisneri M. Bayesian inference for pulsar-timing models. *Mon Not R Astron Soc* 2014, 440:1446–1457.
82. Kipping DM, Nesvornyy D, Buchhave LA, Hartman J, Bakos GA, Schmitt AR. The hunt for exomoons with Kepler (Hek), IV: a search for moons around eight M dwarfs. *Astrophys J* 2014, 784:28–41.
83. Pullen N, Morris RJ. Bayesian model comparison and parameter inference in systems biology using nested sampling. *PLoS One* 2014, 9:e88419.
84. Pártay LB, Bartók AP, Csányi G. Nested sampling for materials: The case of hard spheres. *Phys Rev E* 2014, 89:022302.
85. Pártay LB, Bartók AP, Csányi G. Efficient sampling of atomic configurational spaces. *J Phys Chem B* 2010, 114:10502–10512.
86. Brewer BJ, Pártay LB, Csányi G. Diffusive nested sampling. *Stat Comput* 2011, 21:649–656.
87. Martiniani S, Stevenson JD, Wales DJ, Frenkel D. Superposition enhanced nested sampling. *Phys Rev X* 2014, 4:031034.
88. Fukunishi H, Watanabe O, Takada S. *J Chem Phys* 2002, 116:9058–9067.
89. Mandelshtam VA, Frantsuzov PA, Calvo F. Structural transitions and melting in LJ74-78 Lennard-Jones clusters from adaptive exchange Monte Carlo simulations. *J Phys Chem A* 2006, 110:5326–5332.
90. Nocedal J. Updating quasi-newton matrices with limited storage. *Math Comput* 1980, 35:773–782.
91. Liu DC, Nocedal J. On limited memory BFGS method for large scale optimization. *Math Program* 1989, 45:503–528.
92. Wales DJ. Gmin: A program for basin-hopping global optimisation, basin-sampling, and parallel tempering. Available at: <http://www-wales.ch.cam.ac.uk/software.html>.
93. Haarhoff PC. The density of vibrational energy levels of polyatomic molecules. *Mol Phys* 1963, 7: 101–117.
94. Doye JPK, Miller MA, Wales DJ. Evolution of the potential energy surface with size for lennard-jones clusters. *J Chem Phys* 1999, 111:8417–8428.
95. Steinhardt PJ, Nelson DR, Ronchetti M. *Phys Rev B* 1983, 28:784–805.
96. van Duijneveldt JS, Frenkel D. *J Chem Phys* 1992, 96:4655–4668.