

# Enhancing Torsional Sampling Using Fully Adaptive Simulated Tempering

Miroslav Suruzhon,<sup>1</sup> Khaled Abdel-Maksoud,<sup>1</sup> Michael S. Bodnarchuk,<sup>2</sup> Antonella Ciancetta,<sup>3</sup> Ian D. Wall,<sup>4</sup> and Jonathan W. Essex<sup>1, a)</sup>

<sup>1</sup>*School of Chemistry, University of Southampton, Highfield, Southampton SO17 1BJ, United Kingdom*

<sup>2</sup>*Computational Chemistry, R&D Oncology, AstraZeneca, Cambridge CB4 0WG, United Kingdom*

<sup>3</sup>*Department of Chemical, Pharmaceutical and Agricultural Sciences - DOCPAS, Via Fossato di Mortara 17/19, 44121, Ferrara - University of Ferrara, Italy*

<sup>4</sup>*GSK Medicines Research Centre, Gunnels Wood Road, Stevenage SG1 2NY, United Kingdom*

(Dated: 17 March 2024)

Enhanced sampling algorithms are indispensable when working with highly-disconnected multimodal distributions. An important application of these is the conformational exploration of particular internal degrees of freedom of molecular systems. However, despite the existence of many commonly used enhanced sampling algorithms to explore these internal motions, they often rely on system-dependent parameters, which negatively impact efficiency and reproducibility. Here we present fully adaptive simulated tempering (FAST), a variation of the irreversible simulated tempering algorithm which continuously optimises the number, parameters and weights of intermediate distributions to achieve maximally fast traversal over a space defined by the change in predefined thermodynamic control variable such as temperature or an alchemical smoothing parameter. This work builds on a number of previously published methods, such as sequential Monte Carlo (SMC), and introduces a novel parameter optimisation procedure which can in principle be used in any expanded ensemble algorithm. The method is validated by being applied on a number of different molecular systems with high torsional kinetic barriers. We also consider two different soft-core potentials during the interpolation procedure and compare their performance. We conclude that FAST is a highly efficient algorithm which improves simulation reproducibility and can be successfully used in a variety of settings with the same initial hyperparameters.

Keywords: molecular dynamics | enhanced sampling | simulated tempering | irreversible Monte Carlo | adaptive Monte Carlo

## I. INTRODUCTION

The sampling problem is one of the most significant challenges faced when performing a molecular dynamics (MD) simulation. Many methods that aim to solve this problem have been proposed over the years. One of the most widely used enhanced sampling methods is replica exchange molecular dynamics (REMD)<sup>1</sup>, which uses a series of intermediate distributions to draw samples with increased temperature and diffuse them back into the distribution of interest.

In this work we will focus on the serial and historically much less popular counterpart of REMD—simulated tempering (ST)<sup>2,3</sup>. The main advantage of ST over REMD is that all of the computational time is devoted to a single long simulation, which therefore provides maximum decorrelation of the slow degrees of freedom, such as certain torsions, with respect to their initial values. However, this comes at a price—ST requires the *a priori* unknown relative normalising constants of the intermediate distributions, and is also more sensitive to kinetic trapping in temperature space, since it is no longer guaranteed that every intermediate distribution will produce an equal number of samples. Even though adaptive approaches for estimating the intermediate normalisation constants have been developed<sup>4-7</sup>, ST still remains underutilised compared to REMD<sup>5</sup> arguably due to its lack of robustness.

Both the REMD and ST methods can be generalised beyond temperature to any thermodynamic control variable. The intermediate distributions can represent separate ensembles that

have different values for any specific control variable. In this control variable space, samples can be drawn at regions where traversal of high free energy barriers occurs more frequently and can then be diffused back into the distribution of interest.

The efficiency of tempering algorithms can be further increased by utilising irreversible Markov chain Monte Carlo (MCMC) methods, which are known to provide faster state mixing than their reversible counterparts<sup>8,9</sup>. A straightforward way to make reversible algorithms irreversible is to introduce an additional “lifting coordinate”<sup>10</sup> and enforce an antisymmetric balance condition, known as skew detailed balance<sup>11</sup>. This “lifting coordinate” will control the direction of travel in temperature space explicitly, such that round trips can be enforced. One can then devise a suitable expression for the acceptance criterion which minimises the diffusive motion in temperature space as much as possible<sup>12</sup>. This methodology can be readily applied to a variety of MCMC methods<sup>9</sup>, and irreversible REMD and ST algorithms have been previously published<sup>13-15</sup>. The introduction of this irreversibility can substantially improve the motion in temperature space, which can reach a scaling of  $O(N)$ , where  $N$  is the number of intermediate distributions<sup>13</sup>.

In the following discussion we will present a general algorithm called fully adaptive simulated tempering (FAST) which adaptively determines an optimal parameter protocol for irreversible simulated tempering (IST). While analytical results for selecting optimal parameters exist for reversible REMD<sup>16-18</sup>, irreversible REMD<sup>14,15</sup> and reversible ST<sup>18-20</sup>, a simple analytical expression is not known for IST. Such a methodology is highly desirable not only because of its general applicability, but also because of the well-known fact that a well-performing ST algorithm is more efficient than a well-

<sup>a)</sup>Electronic mail: j.w.essex@soton.ac.uk

performing REMD algorithm<sup>21,22</sup>. One can then use this optimised procedure to maximise the sampling quality of the tempering method.

The following optimisation process has two layers of adaptation: a preliminary offline protocol estimation using adaptive alchemical sequential Monte Carlo (AASMC), as described in<sup>23</sup>, followed by an iterative online protocol refinement during the ST procedure, to increase the number of round trips. The former approach is similar to the methodology considered by Syed *et al.*<sup>14</sup>, while the online optimisation algorithm is the main contribution of this work. In the following work we have focused on round-trip optimisation to maximise the number of transitions between the easiest to sample state and the state of interest.

## II. THEORETICAL BACKGROUND

### A. Simulated Tempering (ST)

ST<sup>2,3</sup> is an expanded ensemble method, which samples from a mixture distribution over the coordinates  $\pi_{mix}(\vec{x})$  which is a weighted sum over each underlying ensemble probability distributions  $\pi_i(\vec{\lambda}_i, \vec{x})$ :

$$\pi_{mix}(\vec{x}) = \sum_{i=1}^{N_\lambda} w_i \pi(\vec{\lambda}_i, \vec{x}) \quad (1)$$

where  $\vec{\lambda}$  is a vector of control variable values (which may or may not be temperatures).  $N_\lambda$  denotes the length of  $\vec{\lambda}$  and so corresponds to the total number of individual thermodynamic states that comprises the expanded ensemble.  $w_i$  is the corresponding weight of the  $i$ -th distribution and  $\pi(\vec{\lambda}_N, \vec{x})$  is the distribution of interest. While the assignment of these weights is arbitrary, they are commonly set to unity, so that all of the distributions are sampled with equal probability. This in turn requires estimation of the normalisation constants for correct weighting.

Exploration in parameter space is commonly done in a Gibbs sampling fashion<sup>24</sup> (Figure 1a), where a state change is attempted after a fixed amount of MD steps. As in REMD, multiple ways of choosing the proposal probabilities  $p_{prop}(\vec{\lambda}_j|\vec{\lambda}_i)$  are possible<sup>24</sup> but a common way of defining the Markov chain is by only attempting transitions between nearest neighbours, with equal proposal probabilities:

$$p_{prop}(\vec{\lambda}_j|\vec{\lambda}_i) = \begin{cases} \frac{1}{2} \delta_{|i-j|,1} & i \notin \{1, N\} \\ \delta_{|i-j|,1} & i \in \{1, N\} \end{cases} \quad (2)$$

with  $\delta$  being the Kronecker delta. The acceptance criterion  $p_{acc}(\vec{\lambda}_j|\vec{\lambda}_i, \vec{x})$  is then related to the importance sampling weight of the configuration  $\vec{x}$ , and is commonly chosen to satisfy detailed balance:

$$p_{acc}(\vec{\lambda}_j|\vec{\lambda}_i, \vec{x}) = \min \left[ 1, \frac{w_j \pi(\vec{\lambda}_j, \vec{x}) p_{prop}(\vec{\lambda}_i|\vec{\lambda}_j)}{w_i \pi(\vec{\lambda}_i, \vec{x}) p_{prop}(\vec{\lambda}_j|\vec{\lambda}_i)} \right] \quad (3)$$




FIG. 1: The Markov chains corresponding to simulated tempering (ST) (Figure 1a) and irreversible simulated tempering (IST) (Figure 1b).

### B. Irreversible Simulated Tempering (IST)

Irreversible simulated tempering (IST)<sup>13</sup> is a generalisation of ST which relaxes the condition of detailed balance (Equation (3)) to that of skew detailed balance<sup>11</sup>. To achieve this, an extra variable  $\sigma$  (“lifting coordinate”) is introduced<sup>10</sup>, thereby creating a mirror image of the irreversible Markov chain of choice (Figure 1b). In this way, the Markov chain increases its state space from  $N_\lambda$  states to  $2N_\lambda - 2$  states. With the skew detailed balance condition, the underlying Markov chain can be designed in any way, as long as all of its states are connected. In IST, the extra variable can be thought of as the direction of the flow—“+1” in the direction  $\vec{\lambda}_1 \rightarrow \vec{\lambda}_{N_\lambda}$  and “-1” in the opposite direction (Figure 1b). This extra variable is strictly +1 at  $\vec{\lambda}_1$ , -1 at  $\vec{\lambda}_{N_\lambda}$  and can take either value at the intermediate  $\vec{\lambda}$  values. This changes the proposal probability  $p_{prop}(\vec{\lambda}_j|\vec{\lambda}_i)$  to:

$$p_{prop}(\vec{\lambda}_j, \sigma_j|\vec{\lambda}_i, \sigma_i) = \begin{cases} \delta_{\sigma_i \sigma_j} & \vec{\lambda}_j = \vec{\lambda}_{i+\sigma_i}, j \notin \{1, N_\lambda\} \\ 1 - \delta_{\sigma_i \sigma_j} & \vec{\lambda}_j = \vec{\lambda}_{i+\sigma_i}, j \in \{1, N_\lambda\} \end{cases} \quad (4)$$

In IST, the skew detailed balance condition enforces the equality:

$$\pi(\vec{\lambda}_i, \sigma_i, \vec{x}) T(\vec{\lambda}_{i+\sigma_i}, \sigma_i|\vec{\lambda}_i, \sigma_i, \vec{x}) = \pi(\vec{\lambda}_{i+\sigma_i}, -\sigma_i, \vec{x}) T(\vec{\lambda}_i, -\sigma_i|\vec{\lambda}_{i+\sigma_i}, -\sigma_i, \vec{x}) \quad (5)$$

with the transition probability  $T(\cdot)$  being the product of the proposal probability  $p_{prop}(\cdot)$  and the acceptance probability  $p_{acc}(\cdot)$ . In this setting, the probability of evolving in  $\vec{\lambda}$  space  $p_{acc}(\vec{\lambda}_j, \sigma_j|\vec{\lambda}_i, \sigma_i, \vec{x})$  is similar to the reversible case (Equation (3)):

$$p_{acc}(\vec{\lambda}_j, \sigma_j|\vec{\lambda}_i, \sigma_i, \vec{x}) = \min \left[ 1, \frac{w_j \pi(\vec{\lambda}_j, \vec{x}) p_{prop}(\vec{\lambda}_i, \sigma_i|\vec{\lambda}_j, \sigma_j)}{w_i \pi(\vec{\lambda}_i, \vec{x}) p_{prop}(\vec{\lambda}_j, \sigma_j|\vec{\lambda}_i, \sigma_i)} \right] \quad (6)$$

where we have omitted the lifting variable from the probability distributions, since it is purely a dummy variable which does not change their functional form. In order to satisfy skew detailed balance, the probability flow in  $\sigma$  space needs to counterbalance the flow in  $\vec{\lambda}$  space. This leads to the lifting coordinate being flipped with probability  $p_{acc}(-\sigma_i|\vec{\lambda}_i, \sigma_i, \vec{x})$ :

$$p_{acc}(-\sigma_i|\vec{\lambda}_i, \sigma_i, \vec{x}) = \max \left[ 0, \sum_{\sigma_k \in \{-1,1\}} T(\vec{\lambda}_{i-\sigma_i}, \sigma_k|\vec{\lambda}_i, -\sigma_i, \vec{x}) - T(\vec{\lambda}_{i+\sigma_i}, \sigma_k|\vec{\lambda}_i, \sigma_i, \vec{x}) \right] \quad (7)$$

The term inside the  $\max[\cdot]$  function is simply the difference between the forward transition probability at the current state and the backward transition probability at the mirrored state. This acceptance criterion is not the only one that satisfies skew detailed balance but is the one that minimises the probability of changing directions<sup>12</sup>. Since the rationale for using IST is precisely the minimisation of diffusive motion in  $\vec{\lambda}$  space, this is the acceptance criterion that will be used hereafter. Finally, the probability of not accepting any  $\vec{\lambda}$  or  $\sigma$  transitions  $p_{rej}(\vec{\lambda}_i, \sigma_i, \vec{x})$  is:

$$p_{rej}(\vec{\lambda}_i, \sigma_i, \vec{x}) = 1 - p_{acc}(\vec{\lambda}_j, \sigma_j|\vec{\lambda}_i, \sigma_i, \vec{x}) - p_{acc}(-\sigma_i|\vec{\lambda}_i, \sigma_i, \vec{x}) \quad (8)$$

In practice, one of  $p_{acc}(\vec{\lambda}_j, \sigma_j|\vec{\lambda}_i, \sigma_i, \vec{x})$ ,  $p_{acc}(-\sigma_i|\vec{\lambda}_i, \sigma_i, \vec{x})$  and  $p_{rej}(\vec{\lambda}_i, \sigma_i, \vec{x})$  is chosen using a pseudo-random number uniformly distributed between 0 and 1, leading to either a transition in  $\vec{\lambda}$  space, a  $\sigma$  flip, or no change. Similarly to ST, the transitions in  $\vec{\lambda}$  and  $\sigma$  space are performed completely independently of the evolution in coordinate ( $\vec{x}$ ) and momentum ( $\vec{p}$ ) space and can in principle be attempted at any point in the simulation.

### C. Multistate Bennett Acceptance Ratio (MBAR)

The multistate Bennett acceptance ratio (MBAR)<sup>25</sup> is a maximum likelihood free energy estimation method, which is also known to be statistically optimal, in the sense of minimising the asymptotic estimator variance. MBAR estimates different thermodynamic observables by creating a self-consistent expanded ensemble model  $\hat{\pi}_{mix}(\vec{\lambda}, \vec{x})$  (here and henceforth, the hat operator denotes an estimated quantity):

$$\begin{aligned} \hat{\pi}_{mix}(\vec{\lambda}, \vec{x}) &= \sum_{i=1}^{N_\lambda} \hat{w}_i \hat{\pi}(\vec{\lambda}_i, \vec{x}) \\ \hat{w}_i &= \frac{N_i}{\sum_{k=1}^{N_\lambda} N_k} \\ \hat{\pi}(\vec{\lambda}_i, \vec{x}) &= e^{-u(\vec{\lambda}_i, \vec{x}) + \hat{f}(\vec{\lambda}_i)} \\ \hat{f}(\vec{\lambda}_i) &= -\ln \left\langle \frac{e^{-u(\vec{\lambda}_i, \vec{x})}}{\hat{\pi}_{mix}(\vec{\lambda}, \vec{x})} \right\rangle_{\hat{\pi}_{mix}(\vec{\lambda})} \end{aligned} \quad (9)$$




FIG. 2: A summary of the FAST workflow.

where  $N_i$  is the total number of samples from  $\vec{\lambda}_i$ ,  $\hat{w}_i$  is the estimated expanded ensemble weight of the  $i$ -th state,  $\hat{f}(\vec{\lambda}_i)$  is the estimated free energy at  $\vec{\lambda}_i$ , usually chosen to be relative to  $\hat{f}(\vec{\lambda}_1)$  and  $u(\vec{\lambda}_i, \vec{x})$  is the dimensionless potential energy. These equations can be solved iteratively using different approaches<sup>25,26</sup>. One can then estimate the expectation value of any observable of interest  $O$  at any intermediate distribution  $\pi(\vec{\lambda}_k)$  using importance sampling, even if this distribution is not explicitly sampled in the observed data:

$$\langle O(\vec{\lambda}, \vec{x}) \rangle_{\pi(\vec{\lambda}_k)} \approx \left\langle O(\vec{\lambda}, \vec{x}) \frac{\hat{\pi}(\vec{\lambda}_k, \vec{x})}{\hat{\pi}_{mix}(\vec{\lambda}, \vec{x})} \right\rangle_{\hat{\pi}_{mix}(\vec{\lambda})} \quad (10)$$

where the expectation  $\langle \cdot \rangle_{\hat{\pi}_{mix}(\vec{\lambda})}$  is obtained by averaging the integrand over each of the total samples. In this work, the samples will be obtained and move proposals in  $\lambda$  space will occur in the same workflow. Samples at the current  $\lambda$  state are collected over  $\tau_{sample}$  units, where  $\tau_{sample}$  is the sampling time between  $\lambda$  change proposals. Afterwards, the resulting MBAR estimator will be used to predict the expectation values of the acceptance criteria described above and estimate transition matrices, as discussed in section Section III A.

## III. DEVELOPING AN ADAPTIVE IST ALGORITHM

An overview of the FAST workflow is provided in Figure 2. The following sections provide detail on each aspect of that workflow.

### A. On-the-Fly Protocol Adaptation

The main contribution of this work is the development of a general adaptive on-the-fly procedure to continuously estimate the optimal protocol  $\vec{\lambda}_{opt} \equiv (\vec{\lambda}_{1,opt}, \dots, \vec{\lambda}_{N_\lambda,opt})$ . We will define  $\vec{\lambda}_{opt}$  to be the protocol which minimises the predicted expected round-trip time  $\hat{\tau}_{round,pred}(\vec{\lambda})$  between  $\vec{\lambda}_1 = 0$  and

$\vec{\lambda}_{N_\lambda} = 1$  (Figure 1b). To obtain this, we first use our MBAR model to estimate an expected transition matrix  $\widehat{\mathbf{T}}(\vec{\lambda})$  connecting all the states in  $\vec{\lambda}$  and  $\sigma$  space. In the case of IST, this translates to a  $(2N_\lambda - 2)$  by  $(2N_\lambda - 2)$  matrix:

$$\widehat{\tau}_{ij}(\vec{\lambda}) \approx \left\langle \widehat{\mathbf{T}}(\vec{\lambda}_j, \sigma_j | \vec{\lambda}_i, \sigma_i, \vec{x}) \frac{\hat{\pi}(\vec{\lambda}_i, \vec{x})}{\hat{\pi}_{mix}(\vec{\lambda}, \vec{x})} \right\rangle_{\hat{\pi}_{mix}(\vec{\lambda})} \quad (11)$$

where each transition probability is calculated using the MBAR free energy estimates according to Equations (6) to (8). In all cases  $\hat{\pi}_{mix}$  is estimated from all previous samples at all  $\vec{\lambda}$  values, as described in Equation (9).

One can then straightforwardly obtain  $\hat{\tau}_{round,pred}(\vec{\lambda})$  by expressing it as the sum of the mean first passage times  $\hat{\tau}_{\lambda=0 \rightarrow \lambda=1} + \hat{\tau}_{\lambda=1 \rightarrow \lambda=0}$ . More generally, the mean first passage time  $\hat{\tau}_{ij}$  from state  $i$  to state  $j$  can be obtained from the equation<sup>27</sup>:

$$\begin{aligned} \hat{\tau}_{ij} &= [(\mathbf{I} - \widehat{\mathbf{T}}_{jj})^{-1} \mathbf{1}]_i \tau_{sample} & i < j \\ \hat{\tau}_{ij} &= [(\mathbf{I} - \widehat{\mathbf{T}}_{jj})^{-1} \mathbf{1}]_{i-1} \tau_{sample} & i > j \end{aligned} \quad (12)$$

where  $\widehat{\mathbf{T}}_{jj}$  is the transition matrix with the  $j$ -th column and row removed,  $\mathbf{I}$  is the identity matrix,  $\mathbf{1}$  is a column vector of ones and  $[\cdot]_i$  denotes the  $i$ -th vector element.

The key assumption behind this methodology is the instantaneous decorrelation of the phase space coordinates—a necessary assumption which is not usually satisfied in real-world applications. Nevertheless, as will be shown later, it is a very useful assumption which works remarkably well in practice, since even dense macromolecular systems often exhibit apparent local energy decorrelation at relatively short timescales (1–10 ps).

The minimisation of  $\hat{\tau}_{round,pred}(\vec{\lambda})$  with respect to  $\vec{\lambda}$  requires an appropriate derivative-free optimisation method. In this work we opt for the covariance matrix adaptation evolution strategy (CMA-ES) algorithm<sup>28</sup>, which is a global optimisation algorithm that is well-known for its robust performance at the number of dimensions relevant to alchemical transformations (typically  $\leq 50$  in practice)<sup>29</sup>. In order to keep the sensitivity of each  $\vec{\lambda}$  value towards variation relatively constant, we will map the  $\vec{\lambda}$  protocol onto an equally-spaced sequence  $\in [0, 1]$  using a piecewise linear interpolation function and perform the optimisation in this transformed space. The optimisation procedure always keeps three  $\vec{\lambda}$  values unchanged: the current  $\vec{\lambda}$  value, 0 and 1. The number of the optimisable  $\vec{\lambda}$  variables will be denoted throughout the text as  $\tilde{N}_\lambda$ .

Once  $\vec{\lambda}_{opt}(\tilde{N}_\lambda)$  corresponding to a particular number of optimisable values  $\tilde{N}_\lambda$  is estimated, the final step is to optimise  $\tilde{N}_\lambda$ . In this work, this will be done using discrete brute-force optimisation, where  $\vec{\lambda}_{opt}(\tilde{N}_\lambda)$  is first calculated at each of  $\min[0, \tilde{N}_\lambda - 1]$ ,  $\tilde{N}_\lambda$  and  $\tilde{N}_\lambda + 1$  dimensions, using the procedure described above. Afterwards,  $\hat{\tau}_{round,pred}(\vec{\lambda}_{opt}(\tilde{N}_\lambda))$  is evaluated at each of these dimensions until a local minimum is found, in which case the minimisation procedure terminates.

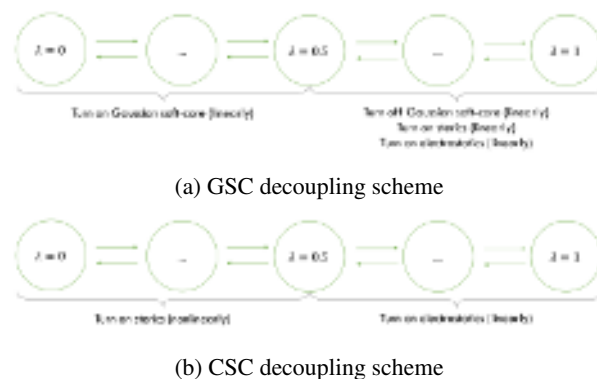


FIG. 3: The main stages of the Gaussian soft-core potential (GSC) (Figure 3a) and the classical soft-core potential (CSC) alchemical schemes (Figure 3b).

In all cases, the initial protocol guess  $\vec{\lambda}_{init}(\tilde{N}_\lambda)$  will either be interpolated in transformed space from the previous minimisation result with the closest (preferably larger) number of dimensions, or taken from the initial adaptive alchemical sequential Monte Carlo (AASMC) run.

Even though the procedure outlined above is theoretically exact and likely to return a globally optimal protocol at infinite sampling, there are still some practical considerations in order for this method to be viable in practice. These will be described below.

## B. Interpolation

The functional form of the alchemical interpolation scheme between the two endpoint distributions  $\pi(0, \vec{x})$  and  $\pi(1, \vec{x})$  does not theoretically influence the sampling at the distribution of interest. Practically, however, the interpolation procedure needs to be carefully chosen to ensure good phase space overlap between all intermediate  $\vec{\lambda}$  windows. To achieve this, soft-core potentials are commonly used for interpolation of Lennard-Jones (LJ) interactions in place of simple linear decoupling. While this results in more efficient sampling, energy evaluations at each  $\vec{\lambda}$  window have to be performed using the whole Hamiltonian and cannot be simply interpolated from the energies at the endpoints. On the other hand, the MBAR estimator uses the evaluated energy of each sample at each previously sampled  $\vec{\lambda}$  value and each iteration of the above protocol optimisation algorithm requires energy evaluations at arbitrary  $\vec{\lambda}$  values chosen by the minimiser as well. Therefore, one would incur unfeasibly high computational expense if full Hamiltonian evaluation was performed throughout these algorithms.

In this work, two different approaches will be compared (Figure 3). The first is always exact and uses a recently developed linearly interpolatable soft-core potential (GSC)<sup>30</sup>, where only three expensive energy evaluations are needed per sample, which can be afterwards stored in memory and readily interpolated as needed. The second approach utilises a commonly used soft-core potential (CSC)<sup>31</sup>, where expensive en-

ergy evaluations will be made at each  $\vec{\lambda}$  value of the protocol. These will then be interpolated linearly in order to approximate intermediate energies for the MBAR estimator and protocol minimiser. Since protocol convergence is asymptotically guaranteed (discussed in Section III E), the energy interpolation error resulting from this procedure at the protocol  $\vec{\lambda}$  values will always tend to zero. This also means that the fraction of samples at the past suboptimal  $\vec{\lambda}$  values diminishes over time, meaning that the GSC approach also results in asymptotically exact free energy values of the converged  $\vec{\lambda}$  protocol, despite being approximate at finite sampling. However, the energy and free energy errors at all non-converged  $\vec{\lambda}$  values will remain finite, meaning that the protocol optimisation procedure is always approximate in this setting and is therefore not guaranteed to converge to the true asymptotically optimal protocol.

### C. Improving MBAR Estimation

Although MBAR is an asymptotically optimal estimator, it is known to produce biased estimates at finite sampling<sup>32</sup>. These can then adversely influence the adaptation process, resulting in local trapping of the free energy and/or protocol estimates and therefore highly suboptimal efficiency. To help tackle this problem, one can use an ensemble of MBAR estimators, using bootstrap aggregation (“bagging”)—a technique often used in machine learning applications to increase the robustness of the estimation<sup>33</sup>. In the current setting, the process consists of simply bootstrapping all available trajectory frames  $N_{boot}$  times and fitting an MBAR estimator to each bootstrapped dataset, resulting in  $N_{boot}$  different, but equally valid, estimators. One can then use the average of the predictions from these models to obtain any observable of interest. In this work, bagging will be used when calculating the Metropolis acceptance criterion and the transition matrix, both of which are dependent on all estimated free energy values  $\vec{f}$  provided by the corresponding MBAR model:

$$\begin{aligned} \hat{p}_{acc}(\vec{\lambda}_j, \sigma_j, \hat{f}_j | \vec{\lambda}_i, \sigma_i, \hat{f}_i, \vec{x}) &\mapsto \langle \hat{p}_{acc}(\vec{\lambda}_j, \sigma_j, \hat{f}_j | \vec{\lambda}_i, \sigma_i, \hat{f}_i, \vec{x}) \rangle_{boot} \\ \hat{\mathbf{T}}(\vec{\lambda}, \vec{f}) &\mapsto \langle \hat{\mathbf{T}}(\vec{\lambda}, \vec{f}) \rangle_{boot} \end{aligned} \quad (13)$$

In order for bootstrapping to generate a correct distribution of the estimators of interest, one needs to supply it with decorrelated samples. Although, as previously discussed, the instant decorrelation assumption is often sufficiently satisfied in practice for timescales on the order of 1–10 ps, it is still essential to obtain a reliable estimate of the effective decorrelation time  $\tau_{decorr}$  for more challenging systems. While a method for estimating  $\tau_{decorr}$  has been previously published<sup>34</sup>, its applicability to FAST is limited due to the constant changes in the  $\vec{\lambda}$  protocol, meaning that an alternative approach needs to be taken.

To obtain an estimate of  $\tau_{decorr}$ , we first note that the round trip time  $\hat{\tau}_{round,pred}(\vec{\lambda})$  predicted by the transition matrix  $\hat{\mathbf{T}}(\vec{\lambda})$  is directly proportional to the sampling time  $\tau_{sample}$  between  $\vec{\lambda}$  proposals (Equation (12)), where it is assumed that  $\tau_{sample}$

provides complete decorrelation at each  $\vec{\lambda}$  value. However, since the transition matrix  $\hat{\mathbf{T}}(\vec{\lambda})$  is independent of  $\tau_{sample}$ , one can also regard the true observed round trip time  $\hat{\tau}_{round,true}$  as being predicted by the same transition matrix  $\hat{\mathbf{T}}(\vec{\lambda})$ , with the only difference being the effective sampling time between  $\vec{\lambda}$  proposals. We now propose that this effective sampling time be equal to an effective decorrelation time  $\tau_{decorr}$ , which can be estimated using the following equation:

$$\hat{\tau}_{decorr} = \frac{\hat{\tau}_{round,true}}{\hat{\tau}_{round,pred}(\vec{\lambda})} \tau_{sample} \quad (14)$$

It should be noted that the ratio  $\frac{\hat{\tau}_{round,true}}{\hat{\tau}_{round,pred}(\vec{\lambda})}$  may not necessarily be independent of  $\tau_{sample}$  in practice. Therefore,  $\hat{\tau}_{decorr}$  is best viewed not as a physical autocorrelation time, but rather as an effective deviation from the instantly decorrelated transition matrix model. In this work, the rounded dimensionless  $N_{decorr} \equiv \max[1, \lfloor \frac{\hat{\tau}_{decorr}}{\tau_{sample}} \rfloor]$  will be used to remove the correlated samples. This will be achieved by starting from an initial sample pseudo-randomly chosen from the most recent  $N_{decorr}$  samples and then keeping only every  $N_{decorr}$ -th previous sample. The resulting effective number of samples  $N_{samples,eff}$  will afterwards be bootstrapped and used for MBAR estimation. Finally, if no round trips have yet been observed,  $\hat{\tau}_{decorr}$  will be estimated from the expected transition time of the longest transition so far observed, instead of the round-trip time.

It is important to note that even though in general different  $\lambda$  states will have significantly different true decorrelation times, it is nevertheless useful to obtain an averaged empirical estimate of them in the form of an effective decorrelation time  $\hat{\tau}_{decorr}$ . This makes both the estimation and the use thereof practically feasible.

### D. Computational Footprint of FAST

Since the computational power required to handle both the estimation of the transition matrix and the free energies scales linearly with respect to the total number of samples  $N_{samples}$ , performing these calculations at a fixed frequency will result in a total computational cost of  $O(N_{samples}^2)$ . To alleviate this, these calculations will in practice be performed at an exponentially diminishing frequency, reducing the complexity to  $O(N_{samples} \log N_{samples})$ . If the implementation is parallelised, written in a compiled language, and/or run in the background on the central processing unit (CPU) while the MD simulation is run on the graphics processing unit (GPU), the computational overhead from adaptation can become negligible with suitably chosen frequency parameters. In this work, the free energies will be calculated every  $\lfloor 1 + 0.01N_{samples,eff} \rfloor$  steps, while protocol optimisation will be performed every  $\lfloor 100 + 0.1N_{samples,eff} \rfloor$  steps, meaning that the number of steps between subsequent optimisations increases linearly with respect to the effective number of samples in both cases.

On the other hand, the memory consumption of the energy matrices required for the MBAR calculations always increases

linearly over time, meaning that depending on the system and the simulation length one might run into memory limitations. In this work, however, matrices with  $\sim 1.6 \times 10^5$  samples (over 160 ns) were routinely handled with memory usage of less than one gigabyte, which is well within the capability of an average computer. Therefore, these considerations are reserved for more computationally intensive cases, where memory requirements could potentially be alleviated by limiting the number of samples used for adaptation, using stochastic approximations of MBAR<sup>35–37</sup>, and/or offloading the matrices to the hard drive using specialised libraries<sup>38</sup>.

The CPU and memory requirements of the free energy estimation procedure are not only dependent on  $N_{\text{samples}}$ , but also on the total  $\vec{\lambda}$  value history, which also grows over time (linearly or logarithmically, depending on the adaptation frequency). This means that exploring a space of continuous (or very high-precision)  $\vec{\lambda}$  values will likely result in unfeasibly high computational requirements and in this work all  $\vec{\lambda}$  values will be preliminarily rounded to two decimal places, meaning that only a maximum of 101  $\vec{\lambda}$  values can be present in the energy matrix. In most practical cases, it is expected that two to three significant figures are completely sufficient for achieving near-optimal performance, while using a relatively low amount of memory.

It is often the case that when one runs MD simulations on GPU-containing compute nodes, they have a large number of CPUs in addition to the GPUs. Not many MD engines can utilise both at the same time, with GROMACS being one of the notable exceptions. OpenMM, on the other hand, does not have this capability, meaning that we can freely run protocol optimization to the unutilised CPU while running exclusively on the GPU, without impacting the simulation performance. Of course, this optimisation would also be possible on the GPU, but then it would interfere with the simulation and decrease its performance in practice. Therefore, an ideal computing architecture is one that has at least a few CPUs in addition to a GPU to aid efficient optimisation. In terms of the parameters discussed above in this section, the only reason we have imposed these limitations is to make the calculations better performing on the compute nodes we used. These parameters are not dependent on the protein-ligand system, but only on the computer hardware specifications.

### E. Convergence

The algorithm described above is highly adaptive and includes the following non-Markovian steps:

- Observable estimation from the MBAR model (free energies and acceptance rates)
- Protocol minimisation
- Estimation of  $\tau_{\text{decorr}}$

Using general results from the literature, it can be shown that the algorithm is asymptotically convergent, despite the multiple layers of adaptation. To demonstrate this, we first state the two sufficient conditions for asymptotic convergence:

containment and diminishing adaptation<sup>39</sup>. The former condition means that if adaptation is stopped at any point, the convergence to the corresponding stationary distribution is guaranteed. This condition is readily satisfied, since all of the above procedures produce finite quantities and non-adaptive IST is still an ergodic sampling algorithm which satisfies skew detailed balance (Equation (5)), even if suboptimal weights and/or  $\vec{\lambda}$  values are used. The second condition can be trivially enforced by using all generated samples for the adaptation, since the variance of any sample-dependent quantity diminishes at infinite sampling. If  $\hat{\tau}_{\text{decorr}}$  is used to remove correlated samples, it also needs to converge to a finite value in order for the above assertion to hold. As can be seen from Equation (14), this convergence is guaranteed as long as the expected round-trip time is finite—a condition which is also satisfied for the IST Markov chain considered hereafter due to its ergodicity (Figure 1b).

### F. Summary of the Method

The full FAST algorithm is shown in Figure 2. The only required input is a set of samples generated at a range of  $\vec{\lambda}$  values with sufficiently good overlap. These can be readily generated by an adaptive sequential Monte Carlo (SMC) algorithm, such the adaptive alchemical sequential Monte Carlo (AASMC) algorithm presented in<sup>23</sup>. Although AASMC requires a number of input parameters, they are all system-independent and only affect the efficiency of the initial stages of the simulation, since FAST eventually converges to an asymptotically optimal protocol regardless of the initial input. In addition, all of the parameters required by the FAST algorithm are related to free energy and/or protocol estimation frequency and quality. Therefore, apart from the degrees of freedom to enhance, all input given to FAST is effectively system-independent, making FAST a nearly black-box enhanced sampling method.

In the following discussion, we will validate FAST on a range of protein-bound and solvated ligand systems, where we will enhance the motions of certain torsional degrees of freedom by decoupling one side of the rotatable bond of interest from the rest of the system at  $\lambda = 0$ . In all cases, the same AASMC and FAST parameters will be used, to demonstrate the wide applicability of the method with minimal prior knowledge.

## IV. RESULTS

### A. Terphenyl in Water

A relatively simple system with an insurmountable kinetic barrier is the terphenyl derivative shown in Figure 4, making it a good test case for alchemical methods. To sample both conformers, one of the 2-tert-butylphenyl groups was completely decoupled at  $\lambda = 0$ . This large alchemical change makes this system one of the more challenging test cases involving a solvated system. FAST simulations of this alchemical change

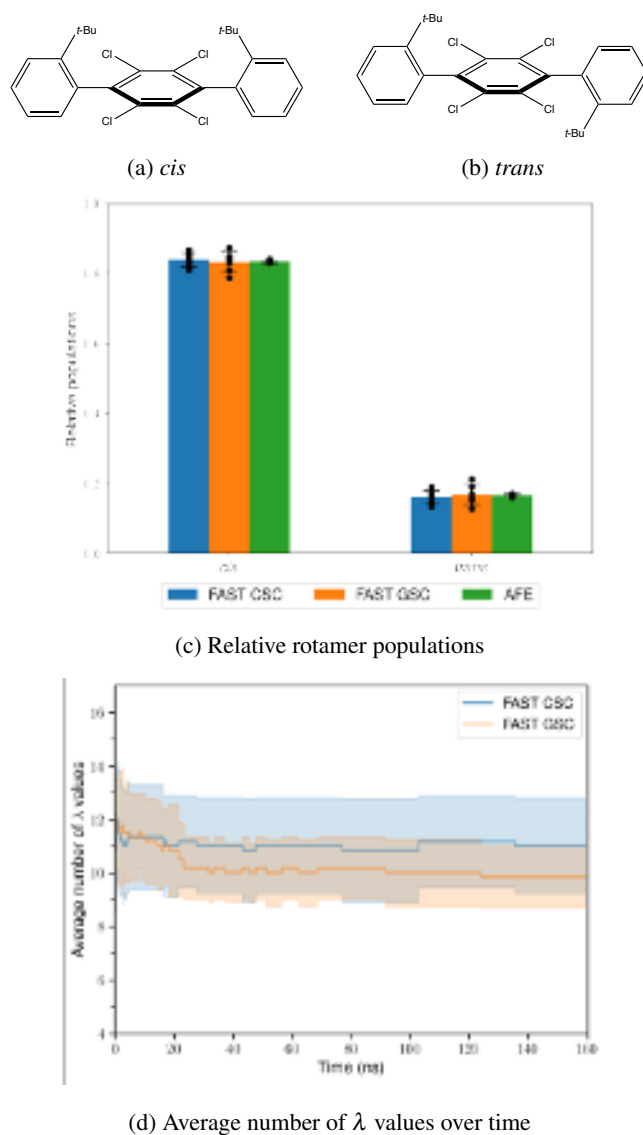


FIG. 4: The two terphenyl derivative rotamers (Figures 4a and 4b), the mean relative populations of these states obtained using FAST and AFE calculations after 6 runs (Figure 4c) and the average number of  $\lambda$  values over time generated during the FAST simulations (Figure 4d). The error bars represent one standard sample deviation.

were validated against alchemical free energy (AFE) calculations. Details of the setup of these and all subsequent AFE calculations can be found in Section 1 of the SI.

The relative rotamer populations, shown in Figure 4c, are well-converged for both the CSC and the GSC soft-core potentials, with the dominant conformer being the *cis* state at  $84\% \pm 2\%$  and  $83\% \pm 3\%$ , respectively. These results are in excellent agreement with the populations obtained from alchemical free energy (AFE) calculations ( $83\% \pm 0\%$ ), indicating that both FAST protocols are able to handle this system without any issues.

As shown in Table S1, both the CSC and the GSC protocols result in a similar number of round trips — approximately 2.5

per nanosecond. This is reflected by the final observed protocol lengths, with an average of  $11 \pm 2$  total  $\lambda$  values for CSC and  $10 \pm 1$  for GSC. These final protocols are approximately 4  $\lambda$  windows shorter on average than the protocol obtained by the initial AASMC run ( $\sim 14$  for both CSC and GSC). This not only showcases the increased efficiency of sampling in  $\vec{\lambda}$  space but also demonstrates the independence of this method on the initial protocol generated by AASMC.

These short final protocol lengths present a somewhat surprising result, because one would expect to need a higher number of intermediate windows for an alchemical perturbation of this size. Indeed, many reported free energy protocols use a higher number of intermediate  $\lambda$  windows for arguably simpler alchemical changes<sup>40–42</sup>. The advantage of using FAST over conventional wisdom is therefore not only the increased robustness and reproducibility of the method compared to manual tuning, but also the increased relative amount of sampling time at  $\lambda = 1$ .

The main potential weakness of ST-based methods is the non-uniform sampling in  $\vec{\lambda}$  space. In this work we will consider the sampling ratio between  $\lambda = 1$  and  $\lambda = 0$ ,  $\frac{N_1}{N_0}$ , and the final relative effective decorrelation time,  $\hat{\tau}_{decorr}$  (Equation (14)), to gauge the sampling reliability of FAST. As shown in Table S2, the CSC protocol results in a less uniform sampling ratio between the two terminal  $\lambda$  values against that observed using GSC. The computed geometric mean sampling ratio between terminal  $\lambda$  values using CSC and GSC was  $\frac{N_1}{N_0} = 1.63 \pm 1.48$  and  $\frac{N_1}{N_0} = 0.97 \pm 1.11$  respectively.  $\hat{\tau}_{decorr}$  is on the other hand comparable between both potentials, with an average value of  $1.37 \pm 0.04$  ps for CSC and  $1.47 \pm 0.07$  ps for GSC (Table S3). Nevertheless, both protocols result in satisfactory sampling ratios and effective decorrelation times, making FAST suitable for the conformational exploration of this system.

## B. T4-lysozyme

Another application of FAST is protein side-chain exploration. One such test case is the Val111 rotation in T4-lysozyme L99A with bound p-xylene (PDB ID: 187L<sup>43</sup>), which has been previously explored with other enhanced sampling methods<sup>44–46</sup> and was also investigated in<sup>23</sup>. Although this is a relatively simple test case, it is a good way to compare the maximum efficiency of both soft-core potentials in this setting.

There are three characteristic conformers for Val111, shown in Figures 5a to 5c. These have previously proven difficult to sample with regular MD<sup>46</sup>, and therefore enhanced sampling methods are needed. In this setting we can achieve this sampling simply by completely alchemically decoupling the Val111 isopropyl group at  $\lambda = 0$  with both FAST and Hamiltonian replica exchange molecular dynamics (H-REMD). Details of the setup of the H-REMD simulations performed for this system and all subsequent systems can be found in Section 1 of the SI.

As shown in Figure 5d, both the CSC and GSC protocols result in statistically equivalent populations after 160 ns of sampling, with the dominant conformer being the *trans* state at

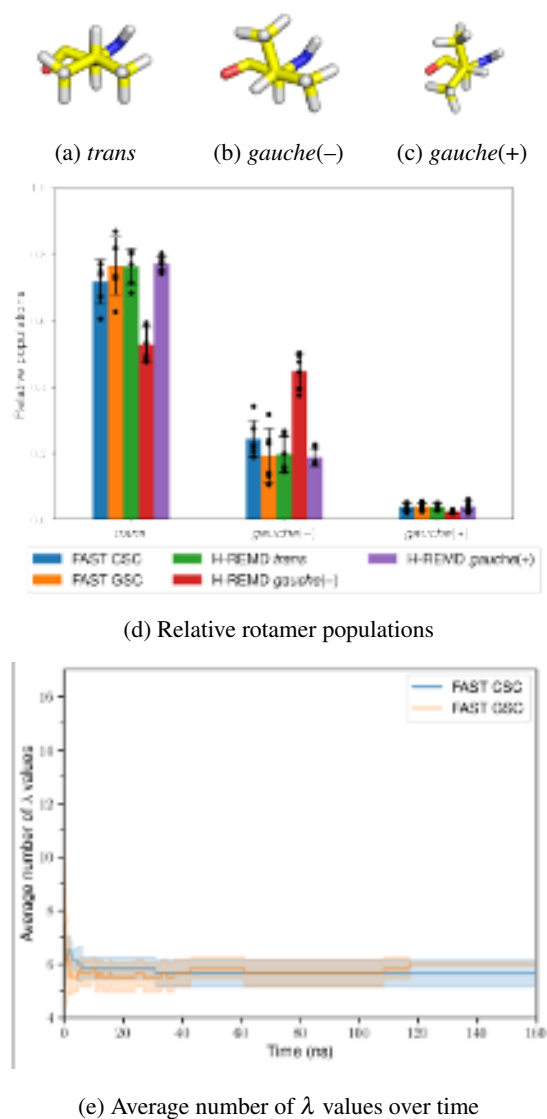


FIG. 5: The three T4-lysozyme Val111 rotamers (Figures 5a to 5c), the mean relative populations of these states obtained using FAST and H-REM after 6 runs (Figure 5d) and the average number of  $\lambda$  values over time generated during the FAST simulations (Figure 5e). The error bars represent one standard deviation.

$72\% \pm 7\%$  and  $77\% \pm 9\%$ , respectively. These results show that the CSC protocol is slightly better converged in this case, similarly to the previous system.

Although both FAST protocols result in variances that are apparently higher than the H-REM ones, the initial H-REM state can significantly affect the final populations even after 160 ns of cumulative sampling over 40  $\lambda$  values. For instance, the *trans* conformer populations are  $76\% \pm 5\%$  and  $53\% \pm 5\%$ , if one starts from the *trans* and *gauche*(-) states, respectively. Therefore, H-REM results in a significant bias towards the initial supplied conformer, meaning that there is insufficient decorrelation from the initial coordinates. This problem is in contrast not observed when FAST is used, since all of the sampling time is dedicated to a single replica. It fol-

lows that despite the lower apparent H-REM variance, there is a higher bias in the resulting population, and if all H-REM simulations are considered together, their cumulative standard deviation becomes 12%, which is higher than either of the FAST protocols.

The resulting round trips per nanosecond are  $6.23 \pm 0.25$  for CSC and  $3.88 \pm 0.38$  for GSC, as shown in Table S1, suggesting that the CSC protocol is significantly more efficient compared to GSC. In comparison, the H-REM protocol results in only  $1.57 \pm 0.08$  round trips per ns (Details of calculation in SI), indicating that the worse-performing FAST protocol is still more than twice as efficient as the unoptimised H-REM protocol. This is evidenced by the low final number of  $\lambda$  values:  $\sim 6$  in both CSC and GSC cases from  $9 \pm 0$  and  $11 \pm 1$  initial  $\lambda$  values respectively (Figure 5e), again showing that the FAST procedure is largely independent of the initial protocol estimated by AASMC.

Similarly to the previous test case, the GSC protocol results in more uniform  $\frac{N_i}{N_0}$  ratios with an average of  $0.93 \pm 1.29$  compared to  $1.59 \pm 1.26$  for CSC (Table S2). However, a higher  $\hat{\tau}_{decorr}$  is observed for GSC:  $2.31 \pm 0.28$  ps compared to  $1.63 \pm 0.05$  ps for CSC (Table S3). It can be therefore concluded that FAST with CSC is more efficient for this test case than FAST with GSC, while H-REM has a very low comparative efficiency to both of the FAST protocols.

### C. Protein Tyrosine Phosphatase 1B (PTP1B)

A practically important use case for enhanced sampling methods is bound ligand conformer sampling. One such test case is PTP1B bound to a thiophene derivative (PDB ID: 2QBS<sup>47</sup>). The rotation of the ligand 3-aminophenyl ring is a rare event whose exploration would be desirable in e.g. binding free energy calculations. This results in two alternative ligand conformers, shown in Figures 6a and 6b. Here we achieve this exploration by completely turning off the 3-aminophenyl ring at  $\lambda = 0$ .

As in the previous test cases, the CSC and GSC protocols result in statistically equivalent populations (Figure 6c):  $90\% \pm 3\% : 10\% \pm 3\%$  and  $87\% \pm 7\% : 13\% \pm 7\%$  with higher variance observed for the GSC protocol. Similar populations are observed for H-REM starting from state 1 and state 2:  $91\% \pm 4\% : 9\% \pm 4\%$  and  $89\% \pm 3\% : 11\% \pm 3\%$ , respectively, showing that in this case H-REM provides results of equivalent quality to the CSC protocol. All of these populations agree with AFE calculations, which result in populations of  $88\% \pm 2\% : 12\% \pm 2\%$ . Similarly to the terphenyl test case, AFE results in lower variance compared to the FAST—an expected behaviour for a low number of conformers. As this number increases, however, AFE calculations become increasingly more impractical. Furthermore, obtaining conformational populations using AFE methods requires prior knowledge of the conformers of interest—knowledge, which is not required by FAST and H-REM.

Both FAST protocols result in a similar number of round trips per nanosecond:  $2.29 \pm 0.35$  for CSC and  $2.24 \pm 0.27$  for GSC (Table S1). These can be compared to  $0.64 \pm 0.08$  for H-REM, meaning that both FAST protocols result in a



This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0190659

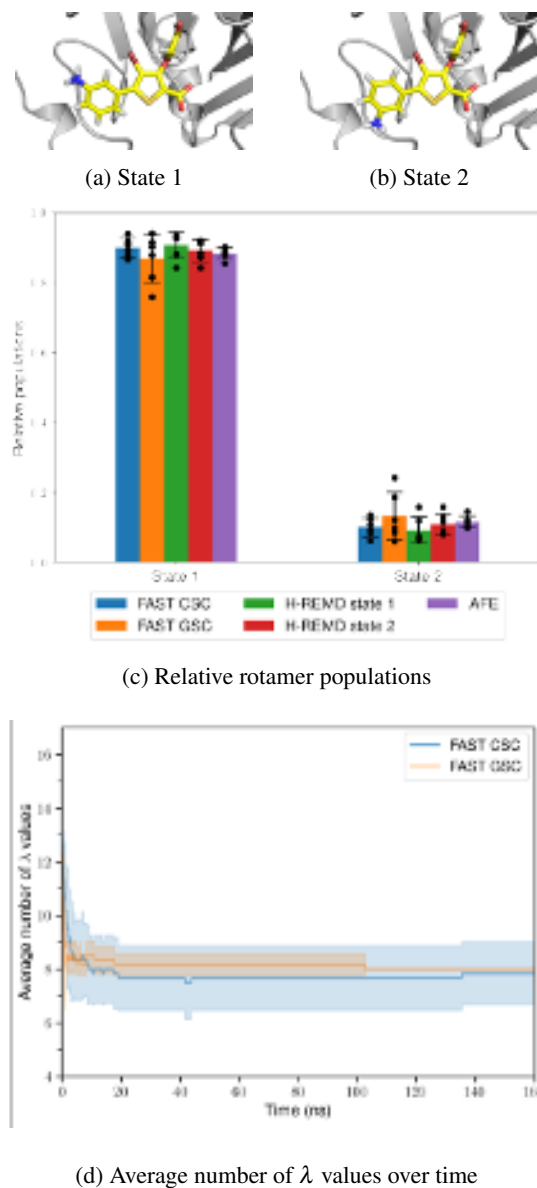


FIG. 6: The two thiophene derivative rotamers bound to PTP1B (Figures 5a to 5c), the mean relative populations of these states obtained using FAST, H-REMD and AFE calculations after 6 runs (Figure 6c) and the average number of  $\lambda$  values over time generated by the FAST simulations (Figure 6d). The error bars represent one standard sample deviation.

nearly fourfold increase in efficiency. This behaviour is again explained by both CSC and GSC resulting in an unexpectedly low total number of  $\lambda$  windows (Figure 6d): 8 in both cases. This presents a significant improvement over the initial 13 (CSC) and 12 (GSC)  $\lambda$  values obtained by AASMC and shows that decoupling a whole phenyl ring does not necessarily require a large number of intermediate steps, as long as the decoupling is performed optimally.

As in the previous systems, the  $\frac{N_1}{N_0}$  ratio is less optimal for the CSC protocol, with a mean value of  $1.36 \pm 1.53$ , compared to  $0.92 \pm 1.32$  for GSC (Table S2). The average  $\hat{\tau}_{decorr}$

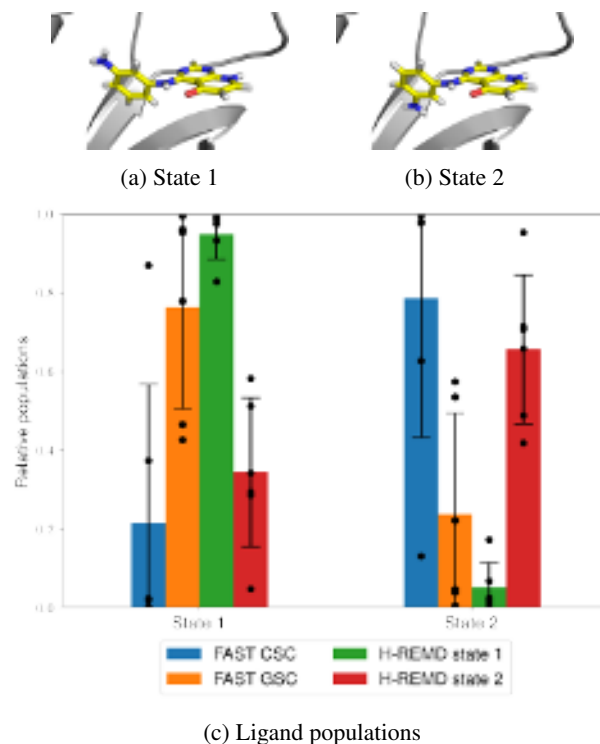


FIG. 7: The two TGF- $\beta$  ligand rotamers (Figures 7a and 7b) and the mean relative populations of these states obtained using the FAST and H-REMD after 6 runs (Figure 7c). The error bars represent one standard sample deviation.

is statistically equivalent in both cases:  $2.10 \pm 0.36$  ps for CSC versus  $2.16 \pm 0.21$  ps for GSC. (Table S3). Therefore, both protocols have comparable performance for this system, with the CSC protocol resulting in lower population variance and the GSC protocol having more consistent  $\frac{N_1}{N_0}$  ratios.

#### D. Transforming Growth Factor Beta (TGF- $\beta$ )

The final test case combines the exploration of a torsional degree of freedom of a ligand bound to transforming growth factor beta (TGF- $\beta$ ) and the nearby Ser82 rotamers. It is experimentally known (PDB ID: 4X2J<sup>48</sup>) that the 4-aminophenyl group of the ligand occupies two alternative states with approximately equal occupancy (Figures 7a and 7b) and that the Ser82 group has three alternative conformations (Figures 8a to 8c). However, a related PDB structure of a more symmetric 3-aminophenyl ligand derivative (PDB ID: 4X2G<sup>48</sup>) was instead used to keep the procedure consistent with our previous publication<sup>23</sup>. Sampling enhancement was achieved by decoupling both the 3-aminophenyl ligand group and the Ser82 hydroxymethyl group at  $\lambda = 0$ .

Both FAST protocols result in highly variable ligand populations (Figure 7c), with state 1 being occupied at  $21\% \pm 35\%$  using the CSC protocol and at  $76\% \pm 26\%$  using the GSC protocol. This high uncertainty is partially observed in the H-REMD runs, where the simulations starting from state 1 stayed in it  $95\% \pm 6\%$  of the time, compared to  $34\% \pm 19\%$  if

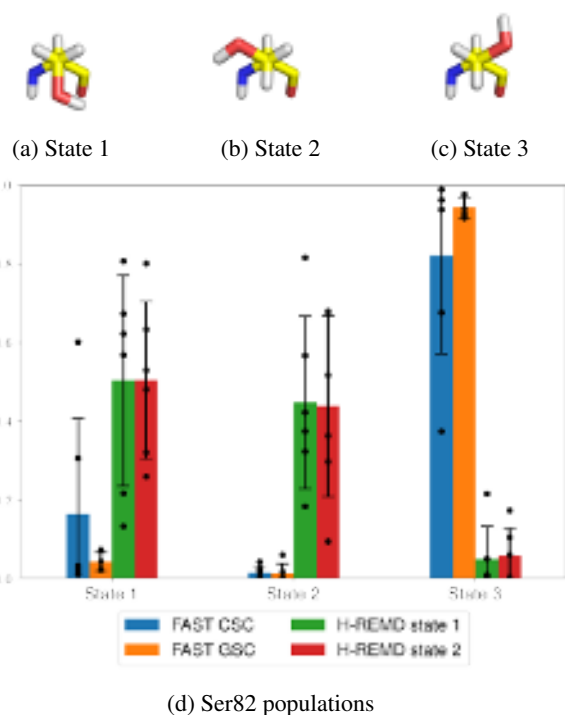


FIG. 8: The three TGF- $\beta$  Ser82 rotamers (Figures 7a and 7b) and the mean relative populations of these states obtained using FAST and H-REMD after 6 runs (Figure 7c). The error bars represent one standard sample deviation.

state 2 is used as an initial state. Similarly to the T4-lysozyme test case, we can see that initial structure biasing is an issue when H-REMD is used and the cumulative variance of all H-REMD results over both initial conformers is 34%, which is comparable to the CSC results.

The Ser82 populations are better converged for GSC than CSC (Figure 8d), with state 3 being occupied at  $94 \pm 3\%$  and  $82\% \pm 25\%$ , respectively. Although these results are consistent between both protocols, they differ significantly from the H-REMD simulations, where a higher variance is observed when both ligand state 1 and state 2 are simulated with state 1 as the initial conformation for Ser82.

These insufficiently converged results can be related to the low number of round trips: only  $0.09 \pm 0.07$  per nanosecond using the CSC protocol compared to  $0.03 \pm 0.02$  with the GSC protocol (Table S1). The H-REMD simulations result in similarly low efficiency with  $0.05 \pm 0.03$  round trips per nanosecond. Since the effective decorrelation time was estimated to be  $50 \pm 25$  ps (CSC) and  $129 \pm 67$  ps (GSC), as shown in (Table S3), it is clear that the instantaneous decorrelation assumption breaks down in this case and kinetic barriers in  $\vec{\lambda}$  space decrease the efficiency of the simulations.

Clustering analysis of the ligand triazanaphthalene ring centre of geometry relative to the initial structure reveals the presence of two main ligand clusters, shown in Figures 9a and 9b. When plotted over time for one of the FAST CSC simulations (Figure 9c), it is revealed that the trappings in  $\vec{\lambda}$  space are correlated with the observed cluster: cluster 1 is more favourable at the lower  $\lambda$  values, while cluster 2 is pre-

ferred at the fully coupled  $\lambda$  values. This behaviour readily explains the low number of round trips, and it can be concluded a slow orthogonal rare event limits the mobility in  $\vec{\lambda}$  space and results in a high effective decorrelation time.

Unfortunately, determining orthogonal rare events *a priori* is not a straightforward task, meaning that systems exhibiting such behaviour are likely to prove challenging for FAST, since the low number of round trips inevitably implies low uniformity of sampling. This is evidenced by the highly inconsistent  $\frac{N_1}{N_0}$  ratios at  $4.91 \pm 5.94$  for CSC and  $6.59 \pm 27.84$  for GSC (Table S2). Nevertheless, the optimisation procedure still results in a relative increase of efficiency, where the initial AASMC  $\lambda$  values are consistently decreased from  $\sim 16$  to 11 for CSC and 10 for GSC (Figure 9d). However, this increase is completely overshadowed by the slow binding mode change.

## V. DISCUSSION

The above results show that FAST is an efficient general-purpose enhanced sampling method of specific internal degrees of freedom, which readily extends the functionality of AASMC to longer timescales. In all of the above test cases, FAST significantly decreases the number of required intermediate states within the initial alchemical protocols provided by AASMC, resulting in a higher proportion of samples being drawn from  $\lambda = 1$ . This results in a better reproducibility of the method, since manual protocol tuning is not required.

However, the choice of the functional form of the interpolation procedure is still a factor which can impact reproducibility. We have shown that although both the CSC and GSC protocols result in similar populations, they exhibit different efficiencies, with GSC consistently resulting in higher dihedral population variance. Moreover, the T4-lysozyme test case demonstrates that GSC can result in a significantly lower round trip rate than CSC. The reason for this is likely suboptimal long-range phase space overlap, which can be explained by the fact that the GSC potential does not always accurately reproduce the real LJ potential, resulting in higher kinetic barriers in  $\vec{\lambda}$  space. This is also evidenced by the PTP1B test case, where the GSC protocol produced populations with higher variance than the H-REMD protocol, even though the number of round trips in the former setting was almost four times higher than the latter.

Interestingly, however, GSC consistently produces more optimal  $\frac{N_1}{N_0}$  ratios (i.e. closer to unity) with lower standard deviations than CSC, where the latter consistently produced samples more highly biased towards  $\lambda = 1$ . This is a surprising result, since with infinite sampling and converged free energy values one would expect these ratios to approach unity and there is no obvious reason why the sampling should be biased in one direction in favour of another. Nevertheless, it appears that GSC is more reliable in this regard, likely due to its often smoother free energy profiles<sup>30</sup>. This smoother profile also presumably leads to better convergence of the  $\vec{\lambda}$  protocol of GSC for T4-lysozyme and PTP1B, as the optimisation problem is closer to a convex one, and there are fewer

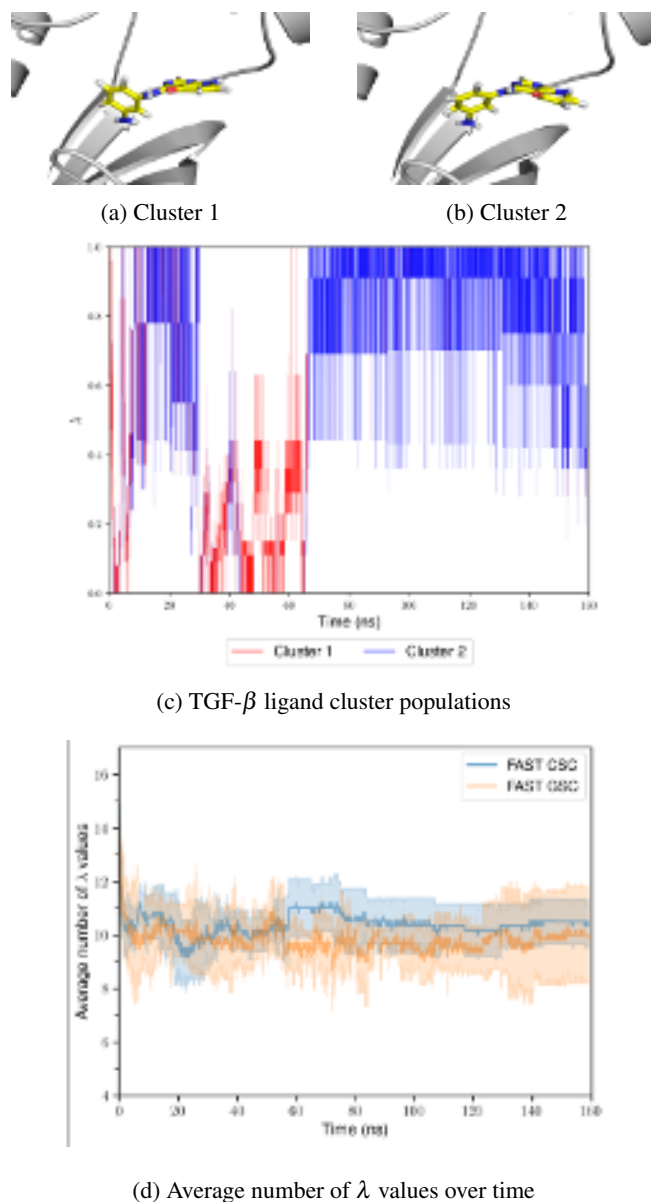


FIG. 9: The two TGF- $\beta$  ligand clusters (Figures 9a and 9b), their transitions over time (Figure 9c) and the average number of  $\lambda$  values over time (Figure 9d).

alternative minima in which the optimiser (CMA-ES) can get stuck relative to CSC.

It has also been demonstrated that  $\tau_{decorr}$  appears to be a very useful metric for determining unexpected kinetic barriers. For instance, it expectedly produces values close to 1 ps in the solvated terphenyl test case, meaning that there is low effective correlation in a homogeneous environment. On the other hand,  $\tau_{decorr}$  is extremely high in the case of TGF- $\beta$ , which immediately hints at orthogonal slow degrees of freedom which impact the sampling negatively. Therefore,  $\tau_{decorr}$  can be monitored in real time to gauge the performance of the FAST sampler if needed.

TGF- $\beta$  is a particularly interesting test case, since it results in a significantly higher variance between different re-

peats compared to the other systems. Even though the nature of the transformation is similar to the other test cases, a substantial increase in  $\tau_{decorr}$  indicates that local exploration of phase space and  $\vec{\lambda}$  space is not as efficient as in the other test cases. As shown in the previous section, this is readily explained by the several alternative binding modes the ligand adopts throughout the simulations. Some of these modes are favourable only in a particular range of  $\lambda$  windows, resulting in significant kinetic trapping and drastic decrease in sampling efficiency. Moreover, any kinetic trapping due to binding modes away from  $\lambda = 1$  indicates that these new modes are not physically relevant and only decrease sampling efficiency to no benefit. This demonstrates the undesirable impact of alchemical decoupling on sampling—it can significantly affect the relative populations in an unexpected way.

Despite the high robustness of FAST on a range of systems, the above test cases show the main weaknesses of the method: unexpected kinetic barriers in  $\vec{\lambda}$  space, as well as slow orthogonal degrees of freedom can significantly affect the sampling efficiency. However, this is a problem which is not unique to FAST, but is more generally relevant to all alchemical/tempering methods using a family of intermediate distributions. Since the slow degrees of freedom are not always known in advance, it will be therefore useful to develop a more general framework to improve long-range phase space overlap either by optimising the functional form of the soft-core potential, or by using e.g. restraint potentials which can help smooth the kinetic barriers in  $\vec{\lambda}$  space. Future work addressing these can therefore help alleviate suboptimal effective decorrelation times.

## VI. CONCLUSION

A fully adaptive version of irreversible simulated tempering has been presented (FAST), where the intermediate distribution protocol is adaptively optimised in real time alongside the relative weights of the distributions. Validation on a variety of systems containing small molecules shows that this method is highly efficient and requires little prior knowledge.

We have also compared two soft-core interpolation methods: classical soft-core potential (CSC) and Gaussian soft-core potential (GSC). In all of the test cases, CSC resulted in lower final distribution variance. Moreover, CSC exhibited higher round trip rates in most cases, as well as lower effective decorrelation times. Nevertheless, more consistent sampling across  $\lambda$  values was observed for GSC, while CSC consistently produced more samples at the distribution of interest ( $\lambda = 1$ ) during the 160 ns of total simulation time.

While the results presented indicate improved sampling efficiency yielded by FAST in comparison to other tempering methods such as H-REMD, the method has been shown to be similarly susceptible to orthogonal slow modes affecting sampling over the alchemical/temperature states. Unforeseen kinetic barriers in the alchemical/temperature space and phase space have been shown to be the main weakness of FAST, and, more generally, alchemical/tempering methods. These have been observed in a protein-ligand system with a slow binding mode transition (TGF- $\beta$ ). Future research will need

to improve the robustness of FAST towards orthogonal slow modes.

## SUPPLEMENTARY MATERIAL

Details regarding the system preparation, simulation protocols, and FAST algorithm performance, are available in the Supporting Information.

## ACKNOWLEDGMENTS

The authors would like to thank Marley L. Samways and Russell Viner for helpful discussions and technical help. This study has been funded by AstraZeneca, GSK and Syngenta, and supported by the EPSRC Centre for Doctoral Training, Theory and Modelling in Chemical Sciences, under Grant EP/L015722/1. This work has also been funded by the EPSRC under Grant EP/V048864/1. The authors acknowledge Iridis 5, the Hartree Centre (JADE, JADE2) and HECBioSim (Grant EP/R029407/1) for the access to high-performance computing clusters which facilitated a part of this study.

## AUTHOR DECLARATIONS

### Conflict of Interest

This work was partly funded by AstraZeneca, Syngenta and GlaxoSmithKline

## DATA AVAILABILITY

All system preparation, the initial AASMC procedure, and the reference Hamiltonian replica exchange molecular dynamics (H-REMD) and alchemical free energy (AFE) calculations have been performed similarly to<sup>23</sup>, while the FAST procedure was performed using OpenMMSLICER 2.0.0<sup>49</sup>. The scripts used to perform all FAST simulations as well as initial structures for such simulations is available to download via a Zenodo repository<sup>50</sup>.

<sup>1</sup>R. H. Swendsen and J.-S. Wang, “Replica Monte Carlo simulation of spin-glasses,” *Physical Review Letters* **57**, 2607–2609 (1986).

<sup>2</sup>E. Marinari and G. Parisi, “Simulated tempering: A new Monte Carlo scheme,” *Europhysics Letters* **19**, 451–458 (1992).

<sup>3</sup>A. P. Lyubartsev, A. A. Martsinovski, S. V. Shevkunov, and P. N. Vorontsov-Velyaminov, “New approach to Monte Carlo calculation of the free energy: Method of expanded ensembles,” *J. Chem. Phys.* **96**, 1776–1783 (1992).

<sup>4</sup>K. M. Åberg, A. P. Lyubartsev, S. P. Jacobsson, and A. Laaksonen, “Determination of solvation free energies by adaptive expanded ensemble molecular dynamics,” *J. Chem. Phys.* **120**, 3770–3776 (2004).

<sup>5</sup>S. Park and V. S. Pande, “Choosing weights for simulated tempering,” *Physical Review E* **76**, 016703 (2007).

<sup>6</sup>P. H. Nguyen, Y. Okamoto, and P. Derreumaux, “Communication: Simulated tempering with fast on-the-fly weight determination,” *J. Chem. Phys.* **138**, 061102 (2013).

<sup>7</sup>Z. Tan, “Optimally adjusted mixture sampling and locally weighted histogram analysis,” *Journal of Computational and Graphical Statistics* **26**, 54–65 (2017).

- <sup>8</sup>P. Diaconis, S. Holmes, and R. M. Neal, “Analysis of a nonreversible Markov chain sampler,” *Ann. Appl. Probab.* **10**, 726–752 (2000).
- <sup>9</sup>F. Faizi, G. Deligiannidis, and E. Rosta, “Efficient irreversible Monte Carlo samplers,” *Journal of Chemical Theory and Computation* **16**, 2124–2138 (2020).
- <sup>10</sup>F. Chen, L. Lovász, and I. Pak, “Lifting markov chains to speed up mixing,” in *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, STOC ’99 (Association for Computing Machinery, New York, 1999) pp. 275–281.
- <sup>11</sup>K. S. Turitsyn, M. Chertkov, and M. Vucelja, “Irreversible Monte Carlo algorithms for efficient sampling,” *Physica D: Nonlinear Phenomena* **240**, 410–414 (2011).
- <sup>12</sup>Y. Sakai and K. Hukushima, “Dynamics of one-dimensional Ising model without detailed balance condition,” *Journal of the Physical Society of Japan* **82**, 064003 (2013).
- <sup>13</sup>Y. Sakai and K. Hukushima, “Irreversible simulated tempering,” *Journal of the Physical Society of Japan* **85**, 104002 (2016).
- <sup>14</sup>S. Syed, A. Bouchard-Côté, G. Deligiannidis, and A. Doucet, “Non-reversible parallel tempering: A scalable highly parallel MCMC scheme,” *J. R. Stat. Soc. Ser. B Stat. Methodol.* **84**, 321–350 (2022).
- <sup>15</sup>S. Syed, V. Romaniello, T. Campbell, and A. Bouchard-Côté, “Parallel tempering on optimized paths,” *arXiv* (2021), 10.48550/ARXIV.2102.07720.
- <sup>16</sup>N. Rathore, M. Chopra, and J. J. de Pablo, “Optimal allocation of replicas in parallel tempering simulations,” *J. Chem. Phys.* **122**, 024111 (2004).
- <sup>17</sup>A. Kone and D. A. Kofke, “Selection of temperature intervals for parallel-tempering simulations,” *J. Chem. Phys.* **122**, 206101 (2005).
- <sup>18</sup>R. Denschlag, M. Lingenheil, and P. Tavan, “Optimal temperature ladders in replica exchange simulations,” *Chemical Physics Letters* **473**, 193–195 (2009).
- <sup>19</sup>F. A. Escobedo and F. J. Martínez-Veracoechea, “Optimized expanded ensembles for simulations involving molecular insertions and deletions. I. Closed systems,” *J. Chem. Phys.* **127**, 174103 (2007).
- <sup>20</sup>F. A. Escobedo and F. J. Martínez-Veracoechea, “Optimization of expanded ensemble methods,” *J. Chem. Phys.* **129**, 154107 (2008).
- <sup>21</sup>S. Park, “Comparison of the serial and parallel algorithms of generalized ensemble simulations: An analytical approach,” *Physical Review E* **77**, 016709 (2008).
- <sup>22</sup>C. Zhang and J. Ma, “Comparison of sampling efficiency between simulated tempering and replica exchange,” *J. Chem. Phys.* **129**, 134112 (2008).
- <sup>23</sup>M. Suruzhon, M. S. Bodnarchuk, A. Ciancetta, I. D. Wall, and J. W. Essex, “Enhancing ligand and protein sampling using sequential Monte Carlo,” *Journal of Chemical Theory and Computation* **18**, 3894–3910 (2022).
- <sup>24</sup>J. D. Chodera and M. R. Shirts, “Replica exchange and expanded ensemble simulations as Gibbs sampling: Simple improvements for enhanced mixing,” *J. Chem. Phys.* **135**, 194110 (2011).
- <sup>25</sup>M. R. Shirts and J. D. Chodera, “Statistically optimal analysis of samples from multiple equilibrium states,” *J. Chem. Phys.* **129**, 124105 (2008).
- <sup>26</sup>X. Ding, J. Z. Vilseck, and C. L. Brooks III, “Fast solver for large scale multistate Bennett acceptance ratio equations,” *Journal of Chemical Theory and Computation* **15**, 799–802 (2019).
- <sup>27</sup>J. Breen and S. Kirkland, “Minimising the largest mean first passage time of a Markov chain: The influence of directed graphs,” *Linear Algebra Appl.* **520**, 306–334 (2017).
- <sup>28</sup>N. Hansen, “The CMA evolution strategy: A comparing review,” in *Towards a New Evolutionary Computation: Advances in the Estimation of Distribution Algorithms*, edited by J. A. Lozano, P. Larrañaga, I. Inza, and E. Bengoetxea (Springer, Berlin, Heidelberg, 2006) pp. 75–102.
- <sup>29</sup>N. Hansen, A. Auger, R. Ros, S. Finck, and P. Pošík, “Comparing results of 31 algorithms from the black-box optimization benchmarking BBOB-2009,” in *Proceedings of the 12th Annual Conference Companion on Genetic and Evolutionary Computation*, GECCO ’10 (Association for Computing Machinery, New York, 2010) pp. 1689–1696.
- <sup>30</sup>Y. Li and K. Nam, “Repulsive soft-core potentials for efficient alchemical free energy calculations,” *Journal of Chemical Theory and Computation* **16**, 4776–4789 (2020).
- <sup>31</sup>T. T. Pham and M. R. Shirts, “Identifying low variance pathways for free energy calculations of molecular transformations in solution phase,” *J. Chem. Phys.* **135**, 034114 (2011).
- <sup>32</sup>J. S. T. Wong, J. J. Forster, and P. W. F. Smith, “Properties of the bridge sampler with a focus on splitting the MCMC sample,” *Statistics and Computing* **30**, 799–816 (2020).

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0190659

- <sup>33</sup>L. Breiman, “Bagging predictors,” *Machine Learning* **24**, 123–140 (1996).
- <sup>34</sup>J. D. Chodera, W. C. Swope, J. W. Pitera, C. Seok, and K. A. Dill, “Use of the weighted histogram analysis method for the analysis of simulated and parallel tempering simulations,” *Journal of Chemical Theory and Computation* **3**, 26–41 (2007).
- <sup>35</sup>B. W. Zhang, J. Xia, Z. Tan, and R. M. Levy, “A stochastic solution to the unbinned WHAM equations,” *J. Phys. Chem. Lett* **6**, 3834–3840 (2015).
- <sup>36</sup>Z. Tan, J. Xia, B. W. Zhang, and R. M. Levy, “Locally weighted histogram analysis and stochastic solution for large-scale multi-state free energy estimation,” *J. Chem. Phys.* **144**, 034107 (2016).
- <sup>37</sup>B. W. Zhang, N. Deng, Z. Tan, and R. M. Levy, “Stratified UWHAM and its stochastic approximation for multicanonical simulations which are far from equilibrium,” *Journal of Chemical Theory and Computation* **13**, 4660–4674 (2017).
- <sup>38</sup>M. Rocklin, “Dask: Parallel computation with blocked algorithms and task scheduling,” in *Proceedings of the 14th Python in Science Conference*, edited by K. Huff and J. Bergstra (SciPy, Austin, 2015) pp. 130–136.
- <sup>39</sup>G. O. Roberts and J. S. Rosenthal, “Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms,” *Journal of Applied Probability* **44**, 458–475 (2007).
- <sup>40</sup>F. Manzoni and U. Ryde, “Assessing the stability of free-energy perturbation calculations by performing variations in the method,” *Journal of Computer-Aided Molecular Design* **32**, 529–536 (2018).
- <sup>41</sup>H. H. Loeffler, S. Bosisio, G. Duarte Ramos Matos, D. Suh, B. Roux, D. L. Mobley, and J. Michel, “Reproducibility of free energy calculations across different molecular simulation software packages,” *Journal of Chemical Theory and Computation* **14**, 5567–5582 (2018).
- <sup>42</sup>T.-S. Lee, B. K. Allen, T. J. Giese, Z. Guo, P. Li, C. Lin, T. D. McGee, D. A. Pearlman, B. K. Radak, Y. Tao, H.-C. Tsai, H. Xu, W. Sherman, and D. M. York, “Alchemical binding free energy calculations in AMBER20: Advances and best practices for drug discovery,” *Journal of Chemical Information and Modeling* **60**, 5595–5623 (2020).
- <sup>43</sup>A. Morton and B. W. Matthews, “Specificity of ligand binding in a buried nonpolar cavity of T4 lysozyme: Linkage of dynamics and structural plasticity,” *Biochemistry* **34**, 8576–8588 (1995).
- <sup>44</sup>L. Wang, B. J. Berne, and R. A. Friesner, “On achieving high accuracy and reliability in the calculation of relative protein–ligand binding affinities,” *Proc. Natl. Acad. Sci. U.S.A.* **109**, 1937 (2012).
- <sup>45</sup>W. Jiang, J. Thirman, S. Jo, and B. Roux, “Reduced free energy perturbation/Hamiltonian replica exchange molecular dynamics method with unbiased alchemical thermodynamic axis,” *J. Phys. Chem. B* **122**, 9435–9442 (2018).
- <sup>46</sup>K. H. Burley, S. C. Gill, N. M. Lim, and D. L. Mobley, “Enhancing side chain rotamer sampling using nonequilibrium candidate Monte Carlo,” *Journal of Chemical Theory and Computation* **15**, 1848–1862 (2019).
- <sup>47</sup>D. P. Wilson, Z.-K. Wan, W.-X. Xu, S. J. Kirincich, B. C. Follows, D. Joseph-McCarthy, K. Foreman, A. Moretto, J. Wu, M. Zhu, E. Binnun, Y.-L. Zhang, M. Tam, D. V. Erbe, J. Tobin, X. Xu, L. Leung, A. Shilling, S. Y. Tam, T. S. Mansour, and J. Lee, “Structure-based optimization of protein tyrosine phosphatase 1B inhibitors: From the active site to the second phosphotyrosine binding site,” *J. Med. Chem.* **50**, 4681–4698 (2007).
- <sup>48</sup>P. Czodrowski, G. Hölzemann, G. Barnickel, H. Greiner, and D. Musil, “Selection of fragments for kinase inhibitor design: Decoration is key,” *J. Med. Chem.* **58**, 457–465 (2015).
- <sup>49</sup>OpenMMSLICER, “<https://github.com/openmmslicer/openmmslicer>,” (2023).
- <sup>50</sup>M. Suruzhon, K. Abdel-Maksoud, M. S. Bodnarchuk, A. Ciancetta, I. D. Wall, and J. W. Essex, “Supporting scripts and data for paper titled “enhancing torsional sampling using fully adaptive simulated tempering - <https://doi.org/10.5281/zenodo.7688833>,” (2023).