

UC Irvine

UC Irvine Previously Published Works

Title

Sampling Conformational Changes of Bound Ligands Using Nonequilibrium Candidate Monte Carlo and Molecular Dynamics.

Permalink

<https://escholarship.org/uc/item/98z2s2ht>

Journal

Journal of Chemical Theory and Computation, 16(3)

ISSN

1549-9618

Authors

Sasmal, Sukanya
Gill, Samuel C
Lim, Nathan M
[et al.](#)

Publication Date

2020-03-10

DOI

10.1021/acs.jctc.9b01066

Peer reviewed



Published in final edited form as:

J Chem Theory Comput. 2020 March 10; 16(3): 1854–1865. doi:10.1021/acs.jctc.9b01066.

Sampling conformational changes of bound ligands using Nonequilibrium Candidate Monte Carlo and Molecular Dynamics

Sukanya Sasmal[†], Samuel C. Gill[‡], Nathan M. Lim[†], David L. Mobley^{†,‡}

[†] Department of Pharmaceutical Sciences, University of California, Irvine

[‡] Department of Chemistry, University of California, Irvine

Abstract

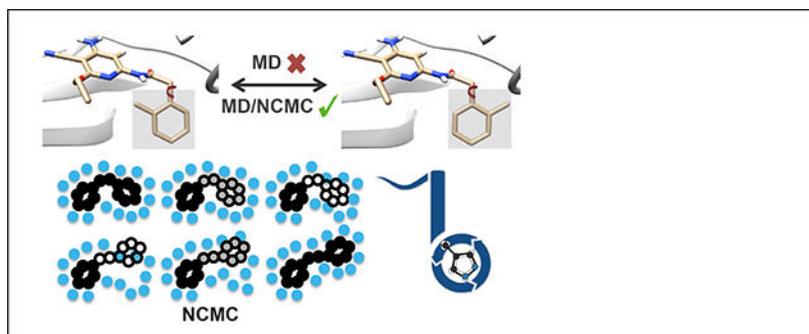
Flexible ligands often have multiple binding modes or bound conformations that differ by rotation of a portion of the molecule around internal rotatable bonds. Knowledge of these binding modes is important for understanding the interactions stabilizing the ligand in the binding pocket, and other studies indicate it is important for calculating accurate binding affinities. In this work, we use a hybrid molecular dynamics (MD)/non-equilibrium candidate Monte Carlo (NMC) method to sample the different binding modes of several flexible ligands and also to estimate the population distribution of the modes. The NMC move proposal is divided into three parts. The flexible part of the ligand is alchemically turned off by decreasing the electrostatics and steric interactions gradually, followed by rotating the rotatable bond by a random angle and then slowly turning the ligand back on to its fully interacting state. The alchemical steps prior to and after the move proposal help the surrounding protein and water atoms in the binding pocket relax around the proposed ligand conformation and increase move acceptance rates. The protein-ligand system is propagated using classical MD in between the NMC proposals. Using this MD/NMC method, we were able to correctly reproduce the different binding modes of inhibitors binding to two kinase targets – c-Jun N-terminal kinase-1 and cyclin-dependent kinase 2 – at a much lower computational cost compared to conventional MD and umbrella sampling. This method is available as a part of the BLUES software package.

Graphical Abstract

dmobley@mobleylab.org.

Supporting Information Available

The Supporting Information is available free of charge on <https://github.com/MobleyLab/bluesflexible-ligand> and includes all the code and scripts used in this work. An archival copy of the version of the tools used in the paper is deposited along with this paper.



1. Introduction

Ligands often bind to proteins in multiple binding modes.^{1,2} Understanding ligand conformational heterogeneity in the binding pocket is important for modulating binding affinity and selectivity, and in turn for having better molecular recognition and function. For example, researchers used knowledge of the two distinct binding modes of mono-substituted ligands of the tyrosine kinase EphB4 to rationally design bi-substituted ligands with single conformation and higher affinity.³ In a different drug design project, nuclear magnetic resonance (NMR) spectroscopy identified minor population states of ligands binding to the antibiotic target LpxC, which hinted towards a larger cryptic binding site and led to the design of a larger and more potent antibiotic.⁴ Thus, improved understanding of binding mechanisms can be used to develop better inhibitors in structure-based drug design.

Multiple binding modes are commonly identified experimentally using X-ray crystallography and NMR techniques, which can be time-consuming or costly. Information about multiple binding modes present in electron density maps is often overlooked during construction of the 3-dimensional structure during X-ray crystallography. One recent study showed that up to 29% of protein-ligand co-crystal structures deposited in the Protein Data Bank have unmodeled ligand conformations and sometimes even unrealistic single bound conformations that are based on averaged electron density contributions from multiple binding modes.¹ The same study also found that less than 2% of ligand-bound crystal structures in the PDB database contained multiple ligand conformations. Additionally, crystal preparation methods like cryocooling can alter the binding modes.^{4,5} Thus, simply having a high-resolution crystal structure does not always guarantee adequate information for structure-guided design.^{4,6} On the other hand, binding affinity studies such as those based on calorimetry present an overall picture of the protein-ligand interactions in solutions, but are hard to interpret in terms of contribution of each individual binding mode. Thus, there is a need for computational tools which can provide information about ligand conformational heterogeneity during rational structure-based drug design.

Moreover, knowledge of ligand binding modes, and in particular the occupancy of each mode, can help improve alchemical binding free energy calculations.^{7,8} These free energy calculations are beginning to be commonly used during lead optimization to prioritize compounds for synthesis and further testing. Ligands with multiple binding modes need separate calculations or other careful study to estimate the contribution of each mode. But, if

the population distribution of the modes is known, then only a single free energy calculation is needed, which decreases the computational cost and improves the accuracy at the same time.

Binding mode inter-conversions can be simple where small rigid fragments flip as in the case of toluene binding to T4-lysozyme⁸ to more complex ones where binding modes differ both in orientation and configuration.^{9,10} Binding mode inter-conversions can be accompanied with a change in the net charge of the ligand,¹¹ which make the systems even more complex. Here we study a binding mode inter-conversion problem of moderate difficulty – where a single rotatable bond in a flexible ligand rotates by 180° to form two distinct modes.

In this work, we introduce a new method for sampling multiple binding modes of flexible ligands with rotatable bonds. We use a hybrid method which combines molecular dynamics (MD) and nonequilibrium candidate Monte Carlo¹² (NCCMC) for this purpose. Previously it has been shown that MD/NCCMC can help to accelerate sampling of multiple binding modes of rigid fragments⁸ and side-chain rotameric states of proteins.¹³

We validated our MD/NCCMC protocol using ligands binding to the protein target c-Jun N-terminal kinase-1 (JNK1). These ligands have different substitutions in a terminal phenyl ring which revolves around a rotatable bond, resulting in two distinct binding modes. Kaus et al. performed extensive alchemical free energy calculations to determine the free energy difference between the two binding modes for these ligands,⁷ which can be converted to population estimates. Additionally, we also validated our method using a different ligand binding to the cyclin-dependent kinase 2 (CDK2) protein with experimental population estimates of the two binding modes based on X-ray electron density maps of crystal structures.¹

Using MD/NCCMC, we computed population estimates close to those obtained using experimental or other computational techniques. MD/NCCMC also required much less computation than standard MD and umbrella sampling, thus suggesting that MD/NCCMC is an efficient computational tool for determining multiple binding modes of flexible ligands.

2. Theory

2.1. Knowledge of binding modes is important for accurate binding free energy calculations

Each binding mode makes a separate contribution to the total binding free energy G° ^{14,15} –

$$\Delta G^\circ = -\beta^{-1} \ln \left(\sum_{i=1}^n e^{-\beta \Delta G_i^\circ} \right)$$

where ΔG_i° is the binding free energy associated with binding mode i , n is the total number of binding modes and β is the inverse of the product of the Boltzmann factor k_B and the absolute temperature T . Subscript i , for the binding mode, refers to only the bound state of the ligand here.

Binding modes often interconvert slowly compared to the timescales of typical MD simulations, and thus interconversions are not generally observed in free energy calculations unless enhanced sampling techniques are used.⁸ As a result, if the binding mode is uncertain or may change as a ligand is modified (e.g. in a lead optimization context) one needs to perform separate free energy calculations limited to each binding mode to compute the total binding free energy.^{7,8} This increases the overall computational cost of binding affinity calculations for ligands with multiple binding modes.

However, if the population distribution of binding modes are known, the the binding affinity calculation can become faster and cheaper. Specifically, if we know the probability p_1 of binding mode 1, then we can use the following two relations to simplify the calculations –

$$p_i \propto e^{-\beta G_{i, bound}^\circ}$$

$$\sum_{i=1} p_i = 1$$

where $G_{i, bound}^\circ$ is the free energy of the binding mode i in the bound form.

The total binding free energy G° can be rewritten as–

$$\Delta G^\circ = \Delta G_1^\circ + \beta^{-1} \ln p_1$$

Thus, only one binding free energy calculation is needed per ligand leading to significant decrease in the computational cost.^{8,16}

2.2. Various computational methods can be applied for sampling binding modes

In this section we mention a few common computational techniques that can be used to sample binding modes, and the benefits and difficulties of using these methods.

The most common and computationally cheapest method for suggesting potential binding modes of ligand is docking.^{17,18} However, docking is designed to rapidly screen large molecular libraries and compromises accuracy to achieve such high speed. As a result, docking is not very efficient in identifying the *true* binding mode(s).¹⁹ Moreover, docking scores do not correlate well with binding affinities¹⁹ and thus, they do not give reasonable estimates of occupancy ratios of multiple binding modes (if present).

An alternative to docking is to use MD-based approaches.²⁰ Classical MD can sample multiple binding modes, but the modes are usually separated by large energy barriers resulting in slower time scales for binding mode transitions in these simulations, necessitating very long simulations. A more pragmatic approach is to use MD-based enhanced sampling techniques to overcome the energy barriers.

In order to overcome energy barriers, metadynamics continuously adapts the potential energy along a few chosen degrees of freedom as the MD simulation progresses which in turn decreases energy barriers between different states. Thus, metadynamics biases the simulation to sample regions of the phase space not yet sampled. This can potentially

enhance the sampling of the entire energy surface. However, metadynamics require some degree of intuition about the important degrees of freedom of the system.

Another way to accelerate sampling across energy barriers is by varying the system temperature in MD simulations, which is commonly done using tempering and annealing methods. For example, replica exchange with solute tempering (REST/REST2)²¹ scales the temperature of only the ligand and selected protein residues to explore alternate conformations. However, there is always a possibility of the protein becoming unstable (and potentially even unfolding) with these techniques because of the high temperatures.

Apart from the works mentioned here, several other binding free energy methods^{22–26} also enhance sampling of internal degrees of freedom and have different strengths and weaknesses.

If the binding modes are known, the problem of crossing energy barriers becomes more tractable. Relative alchemical free energy perturbation (AFE) simulations can sometimes be used in such cases to find out the free energy difference between two binding modes.^{7,27} These calculations simulate a series of separate intermediate states to transform one binding mode into another by disappearing atoms from one region and growing them in another region. The resulting free energy difference can be computed via a variety of perturbation-based estimators such as the multistate Bennett acceptance-ratio method (MBAR), or via alternative formulations like thermodynamic integration (TI). The binding modes need to be very similar to each other for relative AFE calculations.

Another way to determine the free energy difference between binding modes is to use umbrella sampling.²⁸ Similar to AFE, umbrella sampling constructs a series of intermediate states between binding modes along a physical degree of freedom. For example, for sampling the energy change associated with rotating a rotatable bond, the intermediate states are described by different degrees of rotations of the rotatable bond. The system is simulated at the intermediate states using restraints, and then a free energy profile is constructed based on all sampled conformations. However, sampling along a single degree of freedom can result in inadequate sampling of all relevant conformations in complex three-dimensional protein-ligand systems. Particularly, orthogonal degrees of freedom can be especially difficult to adequately sample. Both AFE and umbrella sampling simulations require prior knowledge of what the different binding modes are.

Monte Carlo (MC) is another class of molecular simulation technique where moves are proposed to hop between energy basins. Unlike MD-based techniques, the likelihood of MC moves between binding modes depends only on the initial and proposed states of the system and not on the barrier heights. However, proposing a move without resulting in overlapping atoms is extremely difficult in densely solvated systems. Thus, MD-based techniques are more popular for biomolecular simulations.

Recently, there have been few efforts to use hybrid MD/MC based approaches instead for ligand sampling. For example, Chodera et al. have used rotational and translations MC moves to explore ligand binding modes during absolute binding free energy calculations in the YANK software package.²⁹ Cole et al.³⁰ have combined MC moves involving dihedral

angle rotations with REST to enhance sampling of multiple binding modes during free energy calculations. Also, in our previous work,⁸ MD/MC simulations gave positive results for sampling toluene binding to T4-lysozyme, but the performance degraded with the addition of bulkier side groups (toluene to iodo-toluene). Thus, designing move proposals with high acceptance rates for larger ligands still remains a major challenge.

As discussed in this section, there does not exist a general solution to rapidly sample ligand binding modes and obtain correct populations or free energies, even in cases where binding modes differ primarily only by rotation of one or several rotatable bonds. To address this problem, we developed a hybrid MD and nonequilibrium candidate Monte Carlo based method in this work, which could potentially provide a more general approach.

2.3. Nonequilibrium Candidate Monte Carlo is an efficient sampling tool for crowded systems

Nonequilibrium candidate Monte Carlo (NCCMC) offers the advantages of barrier crossings present in traditional MC moves, but with much higher acceptance rates.¹² Unlike MC where move proposals are *instantaneous*, NCCMC move proposals involve a *finite-time process* where the system is driven through a nonequilibrium protocol to a proposed state.

NCCMC protocol includes a series of perturbation or propagation steps involving chosen degrees of freedom and/or thermodynamic parameters. The switching steps allow the environment surrounding the perturbed region to relax around the proposed state, thereby reducing the chances of steric or electrostatic clashes which are common in MC move proposals. As a result, NCCMC move proposals have higher acceptance rates in densely solvated systems compared to MC.

The proposed NCCMC move is accepted or rejected based on the total work $w[X]$ done during the nonequilibrium process X , given as–

$$w[X] = \sum_{t=1}^T [u_t(x_t) - u_{t-1}(x_t)] + w_{shadow}[X]$$

where x_t is a microstate at switching step t and u_t is the reduced potential energy. The first term in the above equation is the “protocol work” $w_{protocol}$ and the second term is the “shadow work” which accounts for the errors introduced by the use of finite-time-step Langevin integrators.³¹ The acceptance probability $A[X]$ is determined using a modified Metropolis-Hastings criterion³² to maintain detailed balance,

$$A[X] = \min\{1, e^{-w(X)}\}$$

Previously, NCCMC has been used by another group to sample side-chain dihedral angles.³³ In their NCCMC protocol, they used a force to drive the system across energy barriers from one side-chain orientation to another without perturbing any interactions in the system. Although they obtained good results in vacuum systems, the protocol had low acceptance

rates in solvated systems. In particular, the NCMC moves frequently resulted in overlapping water molecules.

2.4. We introduce torsional moves with NCMC for sampling flexible ligands via BLUES software package

In this work, we use the BLUES framework⁸(<https://github.com/MobleyLab/blues>) for implementing NCMC move proposals. The BLUES framework uses a combination of MD and NCMC and is designed specifically for sampling ligand conformations in dense systems.

In BLUES, the NCMC move proposal is defined as a series of alchemical steps where the protein-ligand electrostatic and steric interactions are turned off gradually, followed by proposing a new ligand orientation and then, restoring the interactions. The NCMC switching steps allow the protein and water (if any) surrounding the ligand to relax around the proposed ligand conformation and decreases the chance of overlapping atoms, as observed in the previously described NCMC protocol by Kurut et al.³³

BLUES uses the BAOAB integrator³⁴ for Langevin dynamics, which can sample the configurational space with high accuracy and the shadow work w_{shadow} can be neglected without introducing any significant errors. Thus, the total work w is approximated to be only the protocol work $w_{protocol}$ which is the work done to propagate the system during the NCMC protocol.

In between the NCMC move proposals, BLUES performs conventional MD steps to capture some of the natural protein dynamics by general sampling of the system's motion. Thus, BLUES samples the system using successive NCMC moves and MD propagation steps.

Previously, BLUES has been successfully used to study multiple binding modes of rigid fragments – toluene binding to T4-lysozyme⁸ and caffeine binding to metabolizing enzyme Cytochrome P450 1A2.³⁵ The NCMC protocol used random ligand rotation around the centre of mass of the ligand to sample different binding modes. However, simple rigid-body rotational moves are not sufficient to sample the binding modes of drug-like molecules, which can have conformational degrees of freedom in the binding pocket in addition to rotational and translational ones.

In this work, we extend the NCMC move classes in BLUES by introducing torsional moves for flexible ligands to facilitate transitions between binding modes resulting due to rotation of rotatable bonds. Figure 1 shows the different steps involved in the NCMC protocol. We identify a region connected to the chosen rotatable bond of the ligand which is able to move freely and only this region of the ligand is alchemically turned off and on during the NCMC move proposal. In the middle of the NCMC protocol after the alchemical region is annihilated, the rotatable bond is randomly rotated, followed by restoring the interactions of the alchemical region.

Apart from proposing random torsion rotations of the rotatable bonds in ligands where any amount of rotation (small or large) can be proposed, we also developed another new move

class – flip moves where the rotatable bond is randomly rotated only by large amounts, specifically between 160°–200°.

3. Computational methods

Full details of the methods used in this work can be found in the associated GitHub repository. If we have any updates to the tools/scripts, they will be made available on GitHub. We provide a brief overview of some of the key details of our methods here.

3.1. System preparation

We used four ligands to validate our MD/NCMC protocol (Figure 2). Ligands **1–3** bind to JNK1 and have the same scaffold but with different substitutions in the terminal phenyl group. Their binding free energies were studied extensively by other groups both experimentally³⁶ and computationally.⁷ In the previous computational study with a different force field than that employed here, the ligands exhibited two distinct binding modes with equal probability of occurrence. We modeled the binding mode for these three ligands based on the PDB 2gmx.³⁷ The different aromatic ring substitutions were added to the ligand present in the co-crystal structure using Chimera 1.12.³⁸

For ligand **4**, PDB 2r3i³⁹ was used. Missing residues were modeled using Modeller 9.19.⁴⁰ The partial charges of the ligand atoms were assigned based on the AM1-BCC model⁴¹ using Antechamber, included with the Amber 16 package⁴² and the protonation states of the protein residues using H++⁴³ (<http://biophysics.cs.vt.edu/H++>).

The simulation box was built using tleap, also a part of the Amber 16 package. We used the protein and ligand force field parameters from the Amberff14sb⁴⁴ and GAFF 1.8⁴⁵ force field respectively. Each protein-ligand complex was solvated using TIP3P waters⁴⁶ in a cubic box with 10 Å padding. Two Na⁺ ions and four Cl⁻ ions were added to neutralize the JNK1 and CDK2 protein-ligand complexes respectively.

3.2. Equilibration and production MD simulation details

The MD simulations were performed using OpenMM 7.1.1⁴⁷. We used a Langevin integrator with 2 fs time step and a friction coefficient of 1 ps⁻¹. A cutoff of 10 Å was used for real-space electrostatics and Lennard-Jones forces, while long-range electrostatics was calculated using the particle mesh Ewald method.

Each system was first minimized for 4000 steps with 500 kcal/mol/Å² positional restraints on the heavy atoms present in the protein-ligand complex, followed by another 4000 steps of minimization with 50 kcal/mol/Å² positional restraints only on the protein atoms. Next we performed MD simulations in the NVT ensemble to slowly heat the system and to release the restraints on the protein atoms using the following steps: i) 20 ps of simulation with 25 kcal/mol/Å² restraints at 10 K, ii) gradual heating of the system from 10 to 300 K for 40 ps with 5 kcal/mol/Å² restraints, iii) 20 ps of simulation with 5 kcal/mol/Å² restraints at 300 K, iv) 40 ps of simulation with 1 kcal/mol/Å² restraints, v) 40 ps of simulation with 0.1 kcal/mol/Å² restraints, and vi) 40 ps of simulation without any restraints. In a final phase of

equilibration, we ran the system in the NPT ensemble using a Monte Carlo barostat for 5 ns to get the correct density at 1 atmospheric pressure.

The MD production run was for 100 ns in NVT ensemble with frames being written after every 10 ps. Chimera and CPPTRAJ 16.16⁴⁸ was used for molecular visualizations and MD trajectory analysis respectively.

3.3. MD/NCMC implemented using BLUES software package

We implemented our MD/NCMC protocol using BLUES version 0.1.3. The MD simulations were performed using OpenMM, while the NCMC part used openmmtools 0.14.0⁴⁹ for turning the ligands on and off alchemically. Only the flexible region of the ligand (Figure 2) was selected as the alchemical region. OpenEye Toolkits 3.0.8 (OpenEye Scientific Software) was used to set the selected torsion angles (shown in Figure 2) to random values during NCMC move proposal.

During the NCMC move, all atoms more than 5 Å away from the ligand were frozen. Freezing atoms limits the extent of conformational relaxation to the immediate vicinity of the ligand. This decreases the variance in the βw values and thereby improves the move acceptance rate. This comes at the expense of allowing less long-range relaxation, which has its own costs, but in practice we find that this improves overall acceptance in the systems studied here, as in previous work.⁸

Each iteration consisted of 1000 MD steps and 3400 NCMC switching steps, except in section 4.3 where we varied the number of NCMC switching steps. During the NCMC protocol, the ligand interactions are varied by scaling the λ parameter which control the strength of the non-bonded interactions. The NCMC steps were not evenly distributed across the λ schedule. For λ values between 0.2 to 0.8, we performed additional four perturbation steps for each alchemical step as shown in supplementary Figure S3. For example, for 3400 switching steps, 200 steps were used for λ values 0.0 to 0.2, 3000 steps for λ values 0.2 to 0.8, and 200 steps for λ values 0.8 to 1.0. The NCMC steps were not counted towards the final trajectory. In other words, we only used the MD trajectory as the final MD/NCMC trajectory, thus each iteration corresponded to 1000×0.002 fs/step = 2 ps worth of production run.

We performed a single MD/NCMC simulation for each protein-ligand system and NCMC parameter settings. The uncertainties were calculated based on block-averaging.⁵⁰ We chose the number of blocks in such a way that the standard error of the mean was maximized across the blocks. For probabilities of the modes, we divided the simulation into smaller pieces of 1 ns length for MD/NCMC (10 ns for MD) and then calculated the uncertainties for each chunk of simulation.

3.4. Umbrella sampling performed by sampling different orientations of the rotatable bond

OpenMM was used to carry out the umbrella sampling simulations. We used the same MD simulation parameters described in section 3.2. The rotatable bond of the ligand was harmonically restrained to sample specified dihedral regions in each window. We tried three

different parameter settings for the umbrella sampling simulations -i) 72 windows each with 200 kcal/mol/Å² force constant, ii) 128 windows each with 200 kcal/mol/Å² force constant, and iii) 152 windows with different force constants - 400 or 600 or 800 kcal/mol/Å². Each window was simulated for 10 ns.

While restraining the rotatable bond at different dihedral values in each window, we observed that sometimes it altered the binding mode of the ligand (supplementary Figure S4). Ideally, we would like to sample only a single degree of freedom in umbrella sampling simulations – different orientations of the rotatable bond with the same configuration for the rest of the ligand in our case. Hence, we applied 250 kcal/mol/Å² positional restraint on the fixed part of the ligand (atoms outside the grey box in Figure 2) during simulations with parameter settings (ii) and (iii) to ensure that the protein-ligand interactions are the same for the rest of the ligand.

The one-dimensional free energy profile or the potential of mean force (PMF) was computed using the multistate Bennett acceptance ratio (MBAR) method⁵¹ available through the pymbar package. The population p_i of binding mode i can be calculated using the following equation:

$$p_i = \frac{\sum_{n \in i} e^{-PMF_{w_n}/k_B T}}{\sum_{n \in i, j} e^{-PMF_{w_n}/k_B T}}$$

where w_n corresponds to the window sampling dihedral angle n , PMF_{w_n} is the free energy corresponding to the window w_n . The summation is over windows w_n that collectively sample binding modes i or j .

4. Results

4.1. MD/NCMC is more efficient than MD for sampling the two binding modes

In this work, we developed a hybrid MD/NCMC method to enhance the sampling of binding modes of flexible ligands. We first tested our MD/NCMC method using ligand **1**.

Ligand **1** has a rotatable bond which can rotate even in the bound configuration, giving rise to two different binding modes separated by a 180° rotation of the rotatable bond. We refer these two binding modes as ‘orig’ and ‘flip’ (short for ‘original’ and ‘flipped’, shown in Figure 3). The two phenyl ring orientations are chemically indistinguishable, but distinguishable in simulations because we can number the individual atoms and track which orientation it is in. We should clearly see two distinct and equally populated binding modes in molecular simulations. This is a trivial test but it is a good proof-of-principle because we know what the right answer is and we can test whether we recover this quickly/accurately.

As a point of comparison, we ran a standard 100 ns MD simulation and also monitored the orientation of this ring. Figure 4 shows the binding modes sampled during the MD and MD/NCMC simulations and the probability of the two binding modes as a function of time. We saw only a single transition between the two binding modes during the entire course of the MD simulation, while with MD/NCMC there were frequent transitions between the two

binding modes. Since there were not enough transitions in the MD simulations, the populations of the two binding modes did not converge to the ‘ideal’ 50/50 probability. On the other hand, the populations of the two binding modes converged to the expected values in the MD/NCMC simulations. Thus, MD/NCMC was not only able to increase the transition rate, but it also sampled correctly from the Boltzmann distribution.

It is not straightforward to do a direct efficiency comparison between MD and MD/NCMC simulations since we incur additional computational costs during the NCMC alchemical steps, which are not counted towards the final trajectory length. Force evaluations are the most expensive calculation in simulations and we are interested in sampling more transitions between different binding states. Hence, we compared the two methods based on the observed number of transitions per million force evaluations (transitions/million f-ev).

NCMC requires one force evaluation for each switching step. Thus, for N switching steps and M MD steps, total number of force evaluations performed per iteration is $N + M$. Standard MD executes one force evaluation per step.

Using the calculation described above, we obtain 15.73 ± 0.08 transitions per million f-ev with MD/NCMC, which is about a three order of magnitude improvement over standard MD. We obtain an acceptance rate of $12.7 \pm 0.7\%$ for the NCMC moves, which is similar to 11% obtained by Gill et al.⁸ in previous MD/NCMC work on rigid fragments. Unlike standard Monte Carlo simulations where the ideal efficiency rate is about 20%, the efficiency rate in NCMC simulations is a strong function of the size or the importance of the moves. Particularly even low acceptance of important moves can provide dramatic sampling enhancements, as discussed in the prior NCMC paper.

4.2. Choice of alchemical region important for increasing acceptance rates

During the NCMC move proposal, we preselect only the part of the ligand that moves due to rotation of the rotatable bond. We alchemically turn on and off only this region. However in previous work by Gill et al.,⁸ the whole ligand was specified as the alchemical region. Hence, we investigated the effect of different alchemical regions on move acceptance rates, which in turn affects the convergence time.

We compared the work w distribution for just turning ligand **1** on and off in the binding pocket without any random torsion rotation for two different alchemical regions – i) the whole ligand (Figure 5–a), and ii) only the moving part of the ligand as shown in Figure 2 (Figure 5–b). Ideally we want βw values to be less than or close to zero. This increases the likelihood of moves getting accepted.

When the whole ligand is selected as the alchemical region, the work values are too high (in fact, always significantly unfavorable, here), resulting in every move getting rejected. This is probably because the whole binding pocket undergoes changes when the ligand alchemically disappears (or appears) resulting in high work values. When we designate only the moving part of the ligand as the alchemical region, the work distribution gets shifted to lower values which results in reasonable move acceptance rates.

Thus, in general, we believe that the alchemical region should be selected in such a way that the surrounding protein and water molecules are perturbed only to the extent needed to accommodate the new orientation of the ligand.

It is to be noted that the βw values calculated here are only for turning the ligand on and off in the binding pocket. When a random torsional move is attempted after turning the ligand off in an actual NCMC move proposal, the βw values are overall higher than the case we have looked at in this section since the degree of perturbation in the binding pocket is larger due to rotation of the selected dihedral angle.

4.3. Acceptance rates and actual transitions vary with different amounts of NCMC relaxation

In previous MD/NCMC work,⁸ move acceptance rates were dependent on the NCMC protocols. Hence, we decided to explore how variations in the NCMC protocol, in particular the total amount of NCMC switching steps affect the move acceptance rates (Figure 6).

We see a general trend that with more NCMC relaxation, there is an increase in the acceptance rate. This is expected since more relaxation means more time for the binding pocket to adjust itself around the proposed orientation of the ligand and to relieve any bad contacts.

However, this increase in acceptance rate is deceptive since more NCMC switching steps require more computation. If instead we compare moves accepted per million force evaluations, we see it initially increasing and reaching its maximum value at 3400 NCMC switching steps, after which it decreases. Thus, having more relaxation steps is not always computationally efficient. 3400 switching steps is the optimal amount of relaxation for our current test system.

4.4. Random flip moves increase efficiency of the NCMC protocol

Thus far, we have explored random torsional rotations, which may not be optimal. Particularly, our random torsion rotations can take any value between 0° to 180° . As a result, small rotations are often suggested and accepted, often even within the same binding mode. Not every accepted move thus results in a successful transition between the two binding modes (supplementary Figures S5 and S6).

Move proposals within the same binding mode require additional computation without facilitating transitions between binding modes, thus decreasing the efficiency of MD/NCMC simulations.

An alternative strategy could be to ensure only larger moves are proposed, which could increase efficiency. To test this, we decided to propose only larger moves. Instead of proposing any random value, we designed flip moves, where moves are still randomly chosen but between 160° to 200° values (and thus are suitable only for cases where a torsional angle is expected to have only two stable conformations that differ by roughly 180 degrees).

Figure 7 shows the probability of state as a function of time for MD/NCMC simulation with flip moves for ligand **1**. The moves still sample from Boltzmann distribution and the population converges to the ideal 50/50 population distribution. The number of moves accepted per million force evaluations (30 ± 2) and transitions per million force evaluations (28.4 ± 0.4) are very close in value, whereas for our previous random torsional move proposals, about half of the moves resulted in actual transitions (29 ± 2 moves accepted per million force evaluations and 15.72 ± 0.08 transitions per million force evaluations). The efficiency of flip moves is also evident from the shorter convergence time (~ 2 ns or 1000 iterations) compared to random torsional moves (~ 5 ns or 2500 iterations from Figure 4).

Note that even with flip moves, there are some accepted moves that do not result in an actual transition, as seen from the difference between the number of accepted moves accepted and the number actual transitions. This is probably because during the NCMC relaxation for a small percentage of moves, the ligand is able to relax back to its previous binding mode with acceptable work values.

4.5. MD/NCMC method gives reasonable population estimates for different JNK1 and CDK2 inhibitors

After optimizing and validating our protocol with ligand **1**, we decided to test our MD/NCMC method on additional ligands. Figure 8 shows the probability of the two binding states during the course of MD/NCMC simulations for ligands **2–4**.

Ligands **2** and **3** are from the same chemical series as ligand **1** and differ only by methyl substitutions in the terminal phenyl ring. The populations of the binding modes take different amounts of time to converge for the two ligands – 10 ns or 5000 iterations for ligand **2** and 40 ns or 20000 iterations for ligand **3**. The additional methyl group in ligand **3** decreases the transition rate by almost a factor of seven. Thus, ligands or substituents with bulkier substitutions or large alchemical regions need to be simulated longer to achieve convergence.

Next, we compare our results with population estimates based on AFE calculations by Kaus et al.⁷ For ligand **2**, our population estimates are same as those obtained from AFE calculations (50:50), while for ligand **3**, we get 63:37 population estimates compared to 45:55 obtained using AFE with $\sim 7\%$ error bar. It is to be noted here that the AFE calculations were performed using a different force field (modified OPLS 2.1 for the ligands in the Kaus et al. study), which can give different answers. The final populations are also very sensitive to the AFE protocol. Kaus et al. initially obtained 20:80 population estimates, after which they modified their AFE protocol to obtain a 50:50 estimate for ligand **3**.

MD/NCMC simulations have previously been able to correctly sample binding modes of rigid fragments⁸ and side-chain rotamers.¹³ They also correctly sampled the binding mode populations for ligand **1** in this work. Hence, we believe that the population distribution we obtained using MD/NCMC for ligand **3** is correct given the force fields used in this work.

As a final test, we looked at ligand **4**, which also has two distinct binding modes in CDK2. The occupancies of the two binding modes have been reported to 60:40 based on electron

density maps by Zundert et al.¹ with an error bar of ~10% (occupancies calculated by Zundert et al., and error bars based on personal communication with Henry van den Bedem and James Fraser). This is a computational estimate of the populations based on refinement of the crystal structure. With MD/NCMC, we obtain an occupancy ratio of 53:47, which is within the experimental error margin.

Overall, we find that MD/NCMC significantly improves performance of transition sampling for all three ligands (2–4) which have multiple chemically distinct binding modes. Particularly, for all of these ligands, MD did not allow transitions during the 100 ns long simulations (supplementary Figure S7). The binding modes are likely separated by large energy barriers, which is not easy to overcome with standard MD. Like Monte Carlo moves, the success of NCMC moves is not dependent on the height of the barriers and NCMC moves are able to overcome them. Thus, we see an enhancement in the sampling of binding modes using our MD/NCMC method.

Instead of using direct MD simulations to obtain populations, an alternate approach would have been to employ AFE calculations for each distinct binding mode which, in the best case scenario, could be roughly comparable in cost to MD/NCMC implemented via BLUES. Specifically, standard AFE calculations require separate simulations for 12 lambda values with 5 ns of simulation for each lambda value, resulting in a total 60 ns worth of simulation (based on simulation protocol reported by Kaus et al.⁷). We were able to have sufficiently converged simulations with ~10–20 ns of MD/NCMC simulations, which is computationally equivalent to 44–88 ns of MD simulations. Thus, the computational cost of MD/NCMC simulations is in the same order of magnitude of that for AFE calculations. Note that for n binding modes, we have to perform $n - 1$ AFE calculations to obtain the population distribution, whereas for MD/NCMC we just need to perform a single simulation.

Even though AFE calculations and MD/NCMC may have similar computational cost, at least in the case of few binding modes, AFE calculations are not always that straightforward to implement. For example, the core region of the ligands might require additional restraints to prevent them from drifting out of the binding pocket during the intermediate alchemical states.⁷ Additionally, rotatable bonds might need harmonic restraints such that they sample only a specific binding modes. Furthermore, it is hard to converge AFE calculations to a reasonable accuracy of 1.0 kcal/mol.⁵² In practice, AFE calculations are popularly coupled with tempering methods, such as REST/REST2, for better sampling and convergence. Unlike AFE calculations, with MD/NCMC we have not used any restraints with the ligand having full freedom to move around in the binding pocket if it wants to do so and were able to achieve a satisfactory level of convergence.

4.6. Umbrella sampling requires more computation to converge to the correct population

As an additional point of comparison, we ran umbrella sampling simulations with ligand **1** to compare the computational cost with MD/NCMC simulations. Unlike MD/NCMC simulations, where we were able to get converged simulations easily (or by running longer for larger ligands), we had to optimize and tune the umbrella sampling protocol a number of times to achieve a converged population distribution for the two binding modes (supplementary Figure S8 and S9).

Figure 9 shows the probability of the two states we obtained using different number of windows in umbrella sampling simulations. We had to use 152 windows to achieve convergence. Additional windows directly increases the computational cost. Umbrella sampling was ~35 times more expensive than MD/NCMC simulations for ligand **1**. While we invested considerable effort initially in developing simpler protocols with fewer windows which would perform adequately, we were ultimately unsuccessful — probably in part because steric barriers associated with rotating the group between binding modes often induce ligand and/or protein motions as the bond is rotated.

Additionally, we tried using the protocol optimized with ligand **1** to estimate the population of the binding modes for ligand **2**. We found that same protocol did not work for ligand **2** (supplementary Figure S10). Thus, at least for these binding mode sampling problems, umbrella sampling is not only expensive, it cannot be used as a reliable option for estimating binding modes, especially when doing a prospective or blind study, as the protocol employed would require tuning for each individual ligand to ensure correct populations are obtained — a procedure which is impossible when correct populations are not known.

5. Discussion

In this work, we have developed two NCMC moves – random torsional moves and random flip moves. These two categories of move have different strengths and weaknesses, and the best choice of move likely depends on how much prior information is available. For the ligands studied here, we knew the binding modes and also that the main configuration changes were the flips of the phenyl rings. For such cases, using flip moves is the more efficient choice for the MD/NCMC simulations since the computational effort in NCMC is directed only towards exploring 180° rotations. For systems with limited information, one should perform exploratory simulations with rotational moves to have a better understanding of possible binding modes.

We have studied only multiple binding modes resulting because from ring flips here. However, the torsional moves we designed here can be used for other categories of alternate conformations, for example enhanced sampling of branched components in linear regions of ligands or rotations of terminal rotatable bonds (both are relevant experimentally¹). The NCMC torsional moves can also be combined with other moves like random ligand rotations⁸ or translations to have broader applicability.

Apart from sampling multiple binding modes of flexible ligands, we also tried to simulate flexible ligands in pure water using MD/NCMC to test its efficiency. But, we found out that for such cases the energy barriers separating different conformers could be overcome by standard MD with no need for additional sampling enhancements.

In its current implementation MD/NCMC needs human input about which rotatable bond(s) present in the ligand ought to be explored via this approach. Future work may automate this identifying the rotatable bonds of the ligand and then systematically exploring the binding modes that can arise because of their multiple orientations. For ligands with multiple rotatable bonds, binding modes might differ by multiple torsion angles. For such cases, we

could propose torsional moves of multiple bonds in a single move. We expect that randomly proposing moves for such cases is not an efficient option because the number of possible moves will grow exponentially with increase in the degrees of freedom of the torsional space. To increase acceptance rates, we could draw on exploratory initial simulations of the ligand in implicit or explicit solvent to identify potentially favorable conformations and then propose moves in the binding site randomly based on the explored torsional space.

6. Conclusions

An important component of structure-based drug design is to understand how ligands bind to proteins. However, it is not trivial to determine the binding mode using computational tools. The problem becomes harder when the ligand binds in multiple configurations, or when a new ligand has multiple possible binding modes which have not yet been characterized.

To aid in the problem of binding mode determination and binding mode sampling, here, we developed a new simulation technique based on MD and NCMC to sample binding modes of flexible ligands, and incorporated this into our BLUES package for binding mode sampling. In our approach, we propagate the system using standard MD with regular NCMC move proposals in between to overcome energy barriers. We propose random torsion rotations of rotatable bonds of ligands to accelerate transition rates between different binding modes. However, instead of instantaneous move proposals as in MC, NCMC moves relax the surrounding atoms of the ligand around the proposed orientation to help relax clashes. This helps to improve the move acceptance rate, and in turn the sampling efficiency.

Using MD/NCMC, we were able to sample multiple binding modes of different kinase inhibitors and obtain correct occupancy ratio for the binding modes. The observed transition rates were much faster than standard MD and umbrella sampling. We also found that ‘flip’ moves, where we propose large moves or close to 180° random rotations of rotatable bonds, was about two times more efficient than random torsional moves where any degree of rotation can be proposed. Thus, MD/NCMC has the potential to be a more general and efficient tool for sampling multiple binding modes of flexible ligands.

To sum up, we have applied MD/NCMC to sample binding modes of flexible ligands in this paper and have successfully determined the occupancy ratios of the binding modes of four kinase inhibitors at a much lower computational cost compared to standard MD and umbrella sampling.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgement

We thank Christopher I. Bayly (OpenEye Scientific Software) for valuable scientific discussions on alchemical free energy methods. We also thank Charles L. Brooks III for advising us on umbrella sampling simulations. We gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan Xp used in this research. We also thank OpenEye Scientific Software for providing us with an academic license for OpenEye Toolkits. SS thanks Kalistyn Burley for providing scripts for running and analyzing the umbrella sampling simulations. We also thank Henry van den Bedem and James Fraser for helping us to understand experimental error bars. We also

appreciate financial support from the National Institutes of Health (1R01GM108889-01 and 1R01GM124270-01A1) and the National Science Foundation (CHE 1352608), as well as from OpenEye Scientific Software.

References

- (1). van Zundert GCP; Hudson BM; de Oliveira SHP; Keedy DA; Fonseca R; Heliou A; Suresh P; Borrelli K; Day T; Fraser JS; van den Bedem H qFit-ligand Reveals Widespread Conformational Heterogeneity of Drug-Like Molecules in X-Ray Electron Density Maps. *J. Med. Chem* 2018, 61, 11183–11198. [PubMed: 30457858]
- (2). Mobley DL; Dill KA Binding of small-molecule ligands to proteins: what you see is not always what you get. *Structure* 2009, 17, 489–498. [PubMed: 19368882]
- (3). Bardelle C; Coleman T; Cross D; Davenport S; Kettle JG; Ko EJ; Leach AG; Mortlock A; Read J; Roberts NJ; Robins P; Williams EJ Inhibitors of the tyrosine kinase EphB4. Part 2: structure-based discovery and optimisation of 3, 5-bis substituted anilino-pyrimidines. *Bioorg. Med. Chem. Lett* 2008, 18, 5717–5721. [PubMed: 18851911]
- (4). Lee CJ; Liang X; Wu Q; Najeeb J; Zhao J; Gopalaswamy R; Titecat M; Sebbane F; Lemaitre N; Toone EJ; Zhou P Drug design from the cryptic inhibitor envelope. *Nat. Commun* 2016, 7, 10638. [PubMed: 26912110]
- (5). Fischer M; Shoichet BK; Fraser JS One crystal, two temperatures: Cryocooling penalties alter ligand binding to transient protein sites. *ChemBioChem* 2015, 16, 1560–1564. [PubMed: 26032594]
- (6). Thanos CD; Randal M; Wells JA Potent small-molecule binding to a dynamic hot spot on IL-2. *J. Am. Chem. Soc* 2003, 125, 15280–15281. [PubMed: 14664558]
- (7). Kaus JW; Harder E; Lin T; Abel R; McCammon JA; Wang L How To Deal with Multiple Binding Poses in Alchemical Relative Protein–Ligand Binding Free Energy Calculations. *J. Chem. Theory Comput* 2015, 11, 2670–2679. [PubMed: 26085821]
- (8). Gill SC; Lim NM; Grinaway PB; Rustenburg AS; Fass J; Ross GA; Chodera JD; Mobley DL Binding Modes of Ligands Using Enhanced Sampling (BLUES): Rapid Decorrelation of Ligand Binding Modes via Nonequilibrium Candidate Monte Carlo. *J. Phys.Chem. B* 2018, 122, 5579–5598. [PubMed: 29486559]
- (9). De Simone A; Georgiou C; Ioannidis H; Gupta AA; Juárez-Jiménez J; Doughty-Shenton D; Blackburn EA; Wear MA; Richards JP; Barlow PN; Carragher N; Walkinshaw MD; Hulme AN; Michel J A computationally designed binding mode flip leads to a novel class of potent tri-vector cyclophilin inhibitors. *Chem. Sci* 2019, 10, 542–547. [PubMed: 30746096]
- (10). Hudson L; Mui J; Valázquez S; Carvalho DM; Williams E; Jones C; Bullock AN; Hoelder S Novel Quinazolinone Inhibitors of ALK2 Flip between Alternate Binding Modes: Structure–Activity Relationship, Structural Characterization, Kinase Profiling, and Cellular Proof of Concept. *J. Med. Chem* 2018, 61, 7261–7272. [PubMed: 30085668]
- (11). Schauerl M; Czodrowski P; Fuchs JE; Huber RG; Waldner BJ; Podewitz M; Kramer C; Liedl KR Binding pose flip explained via enthalpic and entropic contributions. *J. Chem. Inf. Model* 2017, 57, 345–354. [PubMed: 28079371]
- (12). Nilmeier JP; Crooks GE; Minh DD; Chodera JD Nonequilibrium candidate Monte Carlo is an efficient tool for equilibrium simulation. *Proc. Natl. Acad. Sci* 2011, 108, E1009–E1018. [PubMed: 22025687]
- (13). Burley KH; Gill SC; Lim NM; Mobley DL Enhancing Side Chain Rotamer Sampling Using Nonequilibrium Candidate Monte Carlo. *J. Chem. Theory Comput* 2019, 15, 1848–1862. [PubMed: 30677291]
- (14). Mobley DL; Chodera JD; Dill KA On the use of orientational restraints and symmetry corrections in alchemical free energy calculations. *J. Chem. Phys* 2006, 125, 084902. [PubMed: 16965052]
- (15). Gallicchio E; Levy RM Recent theoretical and computational advances for modeling protein–ligand binding affinities. *2011*, 85, 27–80.
- (16). Mobley DL; Klimovich PV Perspective: Alchemical Free Energy Calculations for Drug Discovery. *J. Chem. Phys* 2012, 137, 230901. [PubMed: 23267463]

- (17). Huang S-Y; Zou X Advances and Challenges in Protein-Ligand Docking. *Int J Mol Sci* 2010, 11, 3016–3034. [PubMed: 21152288]
- (18). others., et al. D3R Grand Challenge 3: blind prediction of protein–ligand poses and affinity rankings. *J. Comput. Aided. Mol. Des* 2019, 33, 1–18. [PubMed: 30632055]
- (19). Warren GL; Andrews CW; Capelli A-M; Clarke B; LaLonde J; Lambert MH; Lindvall M; Nevins N; Semus SF; Senger S; Tedesco G; Wall ID; Woolven JM; Peishoff CE; Head MS A Critical Assessment of Docking Programs and Scoring Functions. *J. Med. Chem* 2006, 49, 5912–5931. [PubMed: 17004707]
- (20). Williams-Noonan BJ; Yuriev E; Chalmers DK Free energy methods in drug design: Prospects of Alchemical perturbation in medicinal chemistry: miniperspective. *J. Med. Chem* 2017, 61, 638–649. [PubMed: 28745501]
- (21). Wang L; Friesner RA; Berne B Replica exchange with solute scaling: a more efficient version of replica exchange with solute tempering (REST2). *J. Phys. Chem. B* 2011, 115, 9431–9438. [PubMed: 21714551]
- (22). Chia-en AC; Chen W; Gilson MK Ligand configurational entropy and protein binding. *Proc. Natl. Acad. Sci. U.S.A* 2007, 104, 1534–1539. [PubMed: 17242351]
- (23). Jiang W; Roux B Free energy perturbation Hamiltonian replica-exchange molecular dynamics (FEP/H-REMD) for absolute ligand binding free energy calculations. *J. Chem. Theory Comput* 2010, 6, 2559–2565. [PubMed: 21857813]
- (24). Zeifman AA; Stroylov VV; Novikov FN; Stroganov OV; Kulkov V; Chilov GG Alchemical FEP calculations of ligand conformer focusing in explicit solvent. *J. Chem. Theory Comput* 2013, 9, 1093–1102. [PubMed: 26588753]
- (25). Armacost KA; Goh GB; Brooks CL III Biasing potential replica exchange multisite λ -dynamics for efficient free energy calculations. *J. Chem. Theory Comput* 2015, 11, 1267–1277. [PubMed: 26579773]
- (26). Xia J; Flynn W; Levy RM Improving prediction accuracy of binding free energies and poses of HIV integrase complexes using the binding energy distribution analysis method with flattening potentials. *J. Chem. Inf. Model* 2018, 58, 1356–1371. [PubMed: 29927237]
- (27). Michel J; Essex JW Prediction of protein–ligand binding affinity by free energy simulations: assumptions, pitfalls and expectations. *J. Comput. Aided. Mol. Des* 2010, 24, 639–658. [PubMed: 20509041]
- (28). Souaille M; Roux B Extension to the weighted histogram analysis method: combining umbrella sampling with free energy calculations. *Comput. Phys. Commun* 2001, 135, 40–57.
- (29). Chodera JD Yank, <https://github.com/choderalab/yank>.
- (30). Cole DJ; Tirado-Rives J; Jorgensen WL Enhanced Monte Carlo Sampling through Replica Exchange with Solute Tempering. *J. Chem. Theory Comput* 2014, 10, 565–571. [PubMed: 24803853]
- (31). Sivak DA; Chodera JD; Crooks GE Using nonequilibrium fluctuation theorems to understand and correct errors in equilibrium and nonequilibrium simulations of discrete Langevin dynamics. *Phys. Rev. X* 2013, 3, 011007.
- (32). Hastings WK Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 1970, 57, 97–109.
- (33). Kurut A; Fonseca R; Boomsma W Driving Structural Transitions in Molecular Simulations Using the Nonequilibrium Candidate Monte Carlo. *J. Phys. Chem. B* 2017, 122, 1195–1204.
- (34). Leimkuhler B; Matthews C Rational Construction of Stochastic Numerical Methods for Molecular Sampling. *Appl. Math. Res. Express* 2012, 2013, 34–56.
- (35). Jandova Z; Gill SC; Lim NM; Mobley DL; Oostenbrink C Binding modes and metabolism of caffeine. *Chem. Res. Toxicol* 2019, 32, 1374–1383. [PubMed: 31132250]
- (36). Zhao H et al. Discovery of Potent, Highly Selective, and Orally Bioavailable Pyridine Carboxamide c-Jun NH2-Terminal Kinase Inhibitors. *J. Med. Chem* 2006, 49, 4455–4458. [PubMed: 16854050]
- (37). Szczepankiewicz BG et al. Aminopyridine-Based c-Jun N-Terminal Kinase Inhibitors with Cellular Activity and Minimal Cross-Kinase Activity. *J. Med. Chem* 2006, 49, 3563–3580. [PubMed: 16759099]

- (38). Pettersen EF; Goddard TD; Huang CC; Couch GS; Greenblatt DM; Meng EC; Ferrin TE UCSF Chimera – a visualization system for exploratory research and analysis. *J. Comp. Chem* 2004, 25, 1605–1612. [PubMed: 15264254]
- (39). Fischmann TO; Hruza A; Duca JS; Ramanathan L; Mayhood T; Windsor WT; Le HV; Guzi TJ; Dwyer MP; Paruch K; Doll RJ; Lees E; Parry D; Seghezzi W; Madison V Structure-guided discovery of cyclin-dependent kinase inhibitors. *Biopolymers* 89, 372–379. [PubMed: 17937404]
- (40). Eswar N; Webb B; Marti-Renom MA; Madhusudhan M; Eramian D; Shen M.-y.; Pieper U; Sali A Comparative protein structure modeling using Modeller. *Curr. Protoc. Bioinformatics* 2006, 15, 5–6.
- (41). Jakalian A; Jack DB; Bayly CI Fast, efficient generation of high-quality atomic charges. AM1-BCC model: II. Parameterization and validation. *J. Comput. Chem* 2002, 23, 1623–1641. [PubMed: 12395429]
- (42). Case D et al. AMBER 2016, University of California, San Francisco 2016.
- (43). Aguilar B; Anandkrishnan R; Onufriev AV H++ 3.0: automating pK prediction and the preparation of biomolecular structures for atomistic molecular modeling and simulations. *Nucleic Acids Res.* 2012, 40, W537–W541. [PubMed: 22570416]
- (44). Maier JA; Martinez C; Kasavajhala K; Wickstrom L; Hauser KE; Simmerling C ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99SB. *J. Chem. Theory Comput* 2015, 11, 3696–3713. [PubMed: 26574453]
- (45). Wang J; Wolf RM; Caldwell JW; Kollman PA; Case DA Development and testing of a general amber force field. *J. Comp. Chem* 2004, 25, 1157–1174. [PubMed: 15116359]
- (46). Jorgensen WL; Chandrasekhar J; Madura JD; Impey RW; Klein ML Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys* 1983, 79, 926–935.
- (47). Eastman P; Swails J; Chodera JD; McGibbon RT; Zhao Y; Beauchamp KA; Wang L-P; Simmonett AC; Harrigan MP; Stern CD; Wiewiora RP; Brooks BR; Pande VS OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS Comput. Bio* 2017, 13, e1005659. [PubMed: 28746339]
- (48). Roe DR; Cheatham TE PTRAJ and CPPTRAJ: Software for Processing and Analysis of Molecular Dynamics Trajectory Data. *J. Chem. Theory Comput* 2013, 9, 3084–3095. [PubMed: 26583988]
- (49). Chodera J; Rizzi A; Naden L; Beauchamp K; Grinaway P; Fass J; Rustenburg B; Ross GA; Simmonett A choderalab/openmmtools: OpenMMTools 0.13.0. 2017; 10.5281/zenodo.844926.
- (50). Grossfield A; Zuckerman DM Quantifying uncertainty and sampling quality in biomolecular simulations. *Annu. Rep. Comput. Chem* 2009, 5, 23–48. [PubMed: 20454547]
- (51). Shirts MR; Chodera JD Statistically optimal analysis of samples from multiple equilibrium states. *J. Chem. Phys* 2008, 129, 124105. [PubMed: 19045004]
- (52). Wang L; Deng Y; Knight JL; Wu Y; Kim B; Sherman W; Shelley JC; Lin T; Abel R Modeling Local Structural Rearrangements Using FEP/REST: Application to Relative Binding Affinity Predictions of CDK2 Inhibitors. *J. Chem. Theory Comput* 2013, 9, 1282–1293. [PubMed: 26588769]

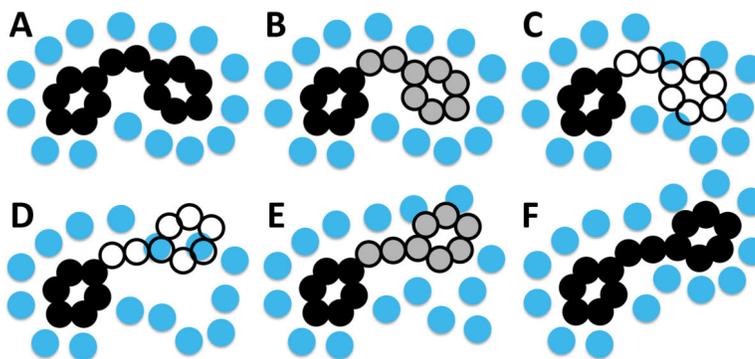


Figure 1.

NCMC move for twisting a rotatable bond of a ligand. Circles with black outline represent the regions of the ligand at various levels of interaction with its neighboring atoms: black – fully interacting, white – non-interacting, and gray – intermediate level. The blue circles are for the surrounding protein and water molecules. A) The ligand is fully interacting. B) The atoms of the ligand which are translocated because of the rotation of the rotatable bond is selected as the alchemical region. The non-bonded interactions of the ligand are slowly decreased only for the alchemical region. This allows the atoms surrounding the alchemical region to relax. C) The alchemical region of the ligand is fully non-interacting. D) The rotatable bond is randomly rotated leading to a new orientation of the alchemical region. E) The alchemical region is slowly turned back on. The surrounding atoms relax around the new configuration of the ligand to relieve clashes if present. F) The whole ligand is fully interacting. The total work done in performing steps A-F is calculated and the NCMC move is accepted or rejected based on it.

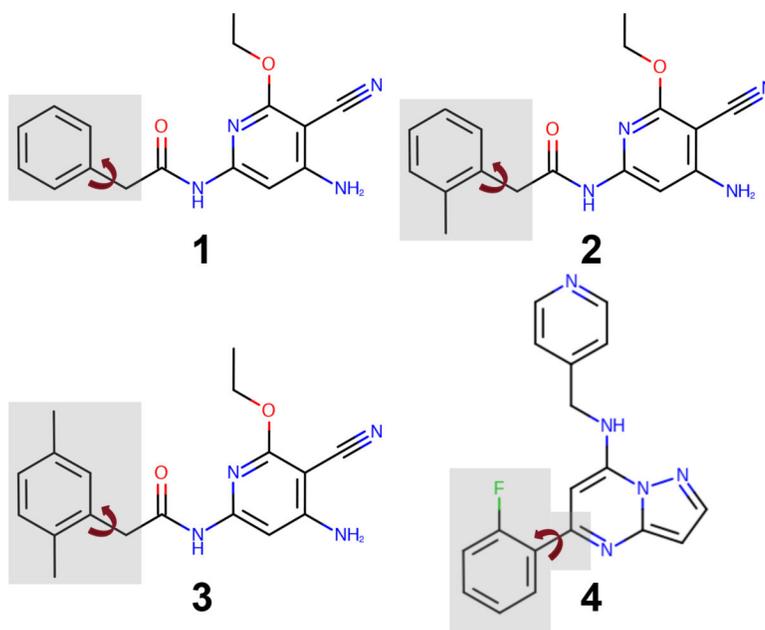


Figure 2. Structure of the four ligands used to validate our MD/NCMC protocol. Ligands **1–3** bind to JNK1, while ligand **4** binds to CDK2. The grey boxes show the regions selected as the alchemical regions for the NCMC moves. The arrows indicate the bonds rotating freely in the bound form of the ligands. The remainder of the ligand forms an anchor or core which retains its pose (key interactions with the protein) in the binding site as the remainder of the ligand rotates. More details about the protein-ligand interactions can be found in the contact maps presented in the supplementary Figures S1 and S2.

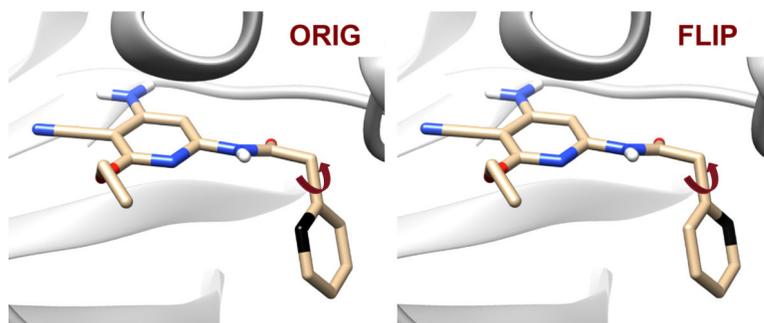


Figure 3. Binding modes of ligand **1**. The rotation of the rotatable bond leads to two distinguishable binding modes in molecular simulations - 'orig' and 'flip'. The black atom in the phenyl ring represents a carbon atom and is colored black for reference only.

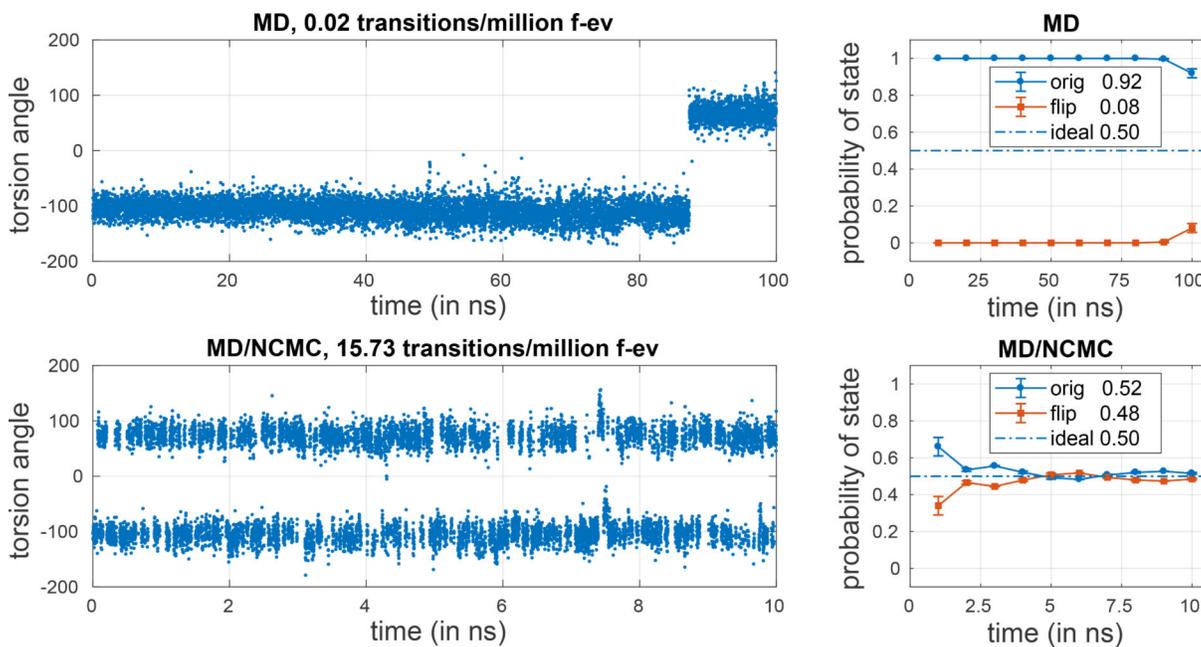


Figure 4.

Comparison of the sampling efficiency of MD and MD/NCMC simulations for ligand **1**.

Torsion angles (corresponding to the two binding modes) sampled during the course of the simulations as a function of time (*left*). The number of transitions between the two binding modes per million force evaluations (transitions/million f-ev) is also reported. Probability of the two binding modes, namely 'orig' and 'flip' as a function of time (*right*). MD/NCMC increases the rate of transition between the two binding modes compared to classical MD simulations, resulting in faster convergence of the population of two binding modes to the 'ideal' distribution.

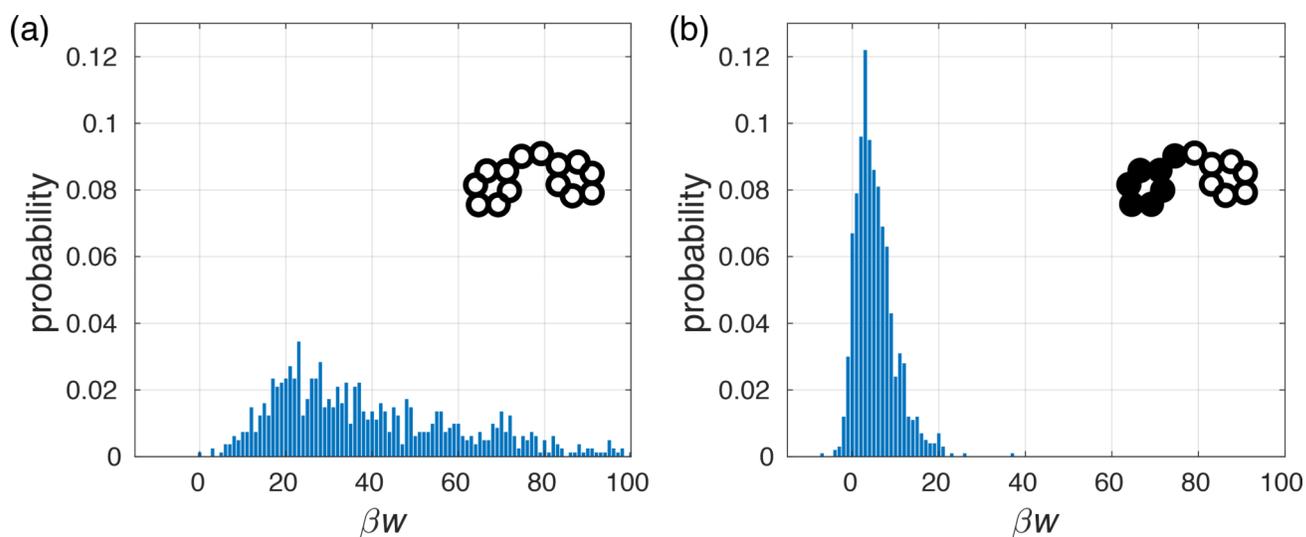


Figure 5. Distribution of work w done in turning ligand **1** on and off in the binding pocket without proposing any random torsion rotation. The alchemical region is specified as follows: a) the whole ligand; b) only the flexible part of the ligand. β is the inverse of the product of the Boltzmann constant k_B and the temperature T . Proposed moves with smaller values of βw (close to or less than zero) are more likely to be accepted. Thus, defining only the flexible part of the ligand as the alchemical region is better for move acceptances as seen from the work distribution.

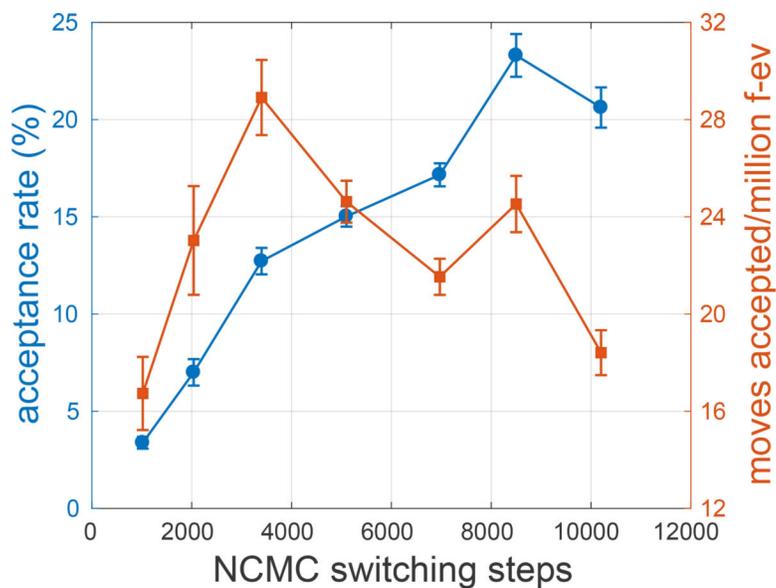


Figure 6. Acceptance rate of NCMC moves (in blue) and moves accepted per million force evaluations (f-ev) (in orange) as a function of the number of NCMC switching steps for ligand **1**. The acceptance rate increases with more NCMC steps. However, the number of moves accepted per million force evaluations initially increases and then decreases with the number of NCMC switching steps.

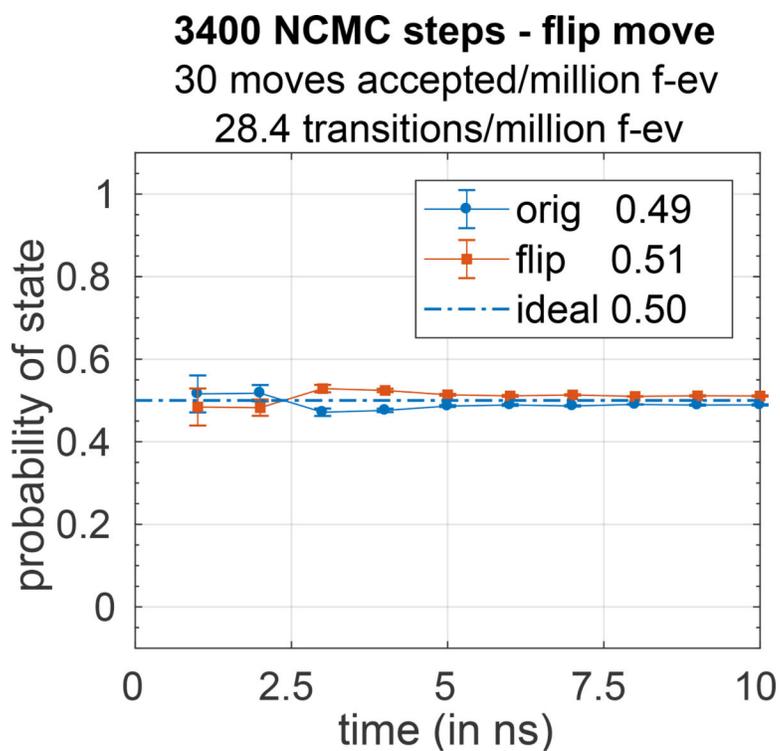


Figure 7. Probability of the two binding states of ligand **1** as a function of time sampled using MD/ NCMC with flip moves. The simulation converges around 2x faster than our previous simulation with random rotational moves in Section 4.1.

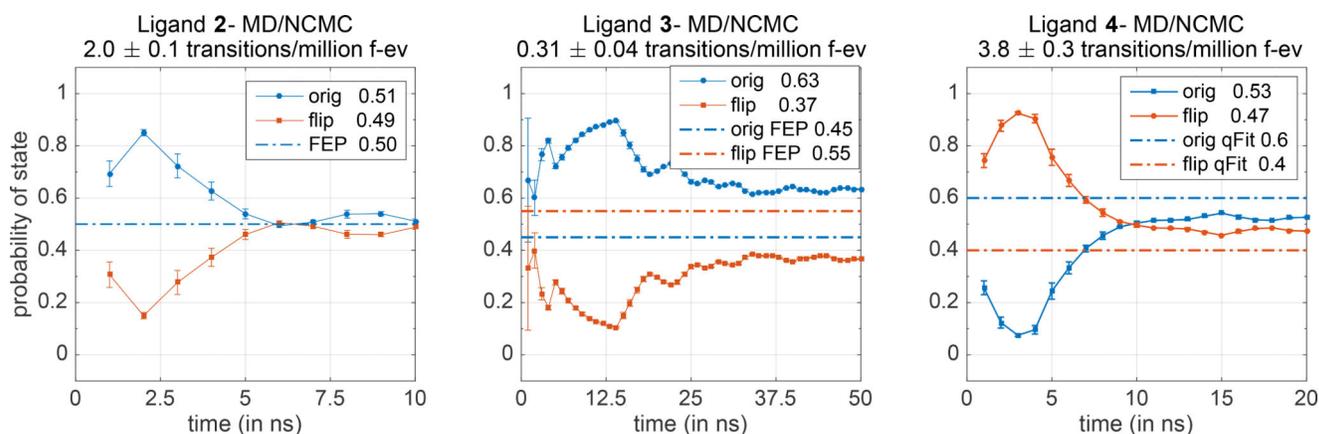


Figure 8.

Probability of the binding states of ligands **2**, **3** and **4** as a function of time based on MD/NCMC simulations with flip moves (solid lines). The simulations require different amounts of time to converge for the three ligands. The final population estimates of the binding modes for ligand **2** is close to that obtained using AFE calculations by Kaus et al.⁷ (dashed lines), whereas for ligand **3**, the converged populations are slightly different from AFE estimates (dashed lines). For ligand **4**, the population estimates are similar to those computed by Zundert et al.¹ using X-ray electron density data with an error margin of 0.1 (dashed lines).

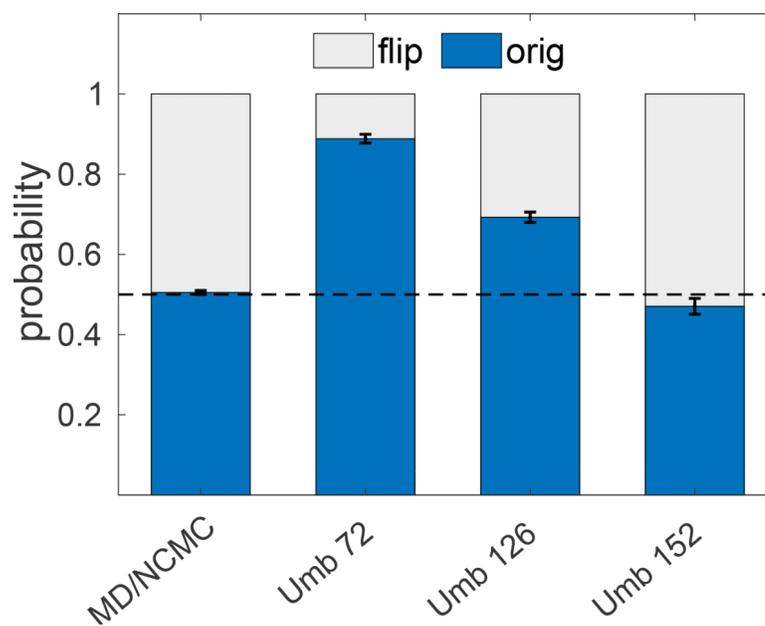


Figure 9. Probability of the two binding modes obtained using MD/NCMC and umbrella sampling (Umb) with different number of windows. The ideal population estimate is 50:50 shown by the dotted line. We had to use 152 windows to get estimates close to the ideal value with umbrella sampling, which increased the computational cost considerably.