

A fast, open source implementation of adaptive biasing potentials uncovers a ligand design strategy for the chromatin regulator BRD4

Bradley M. Dickson,^{1,a)} Parker W. de Waal,² Zachary H. Ramjan,¹ H. Eric Xu,^{2,3} and Scott B. Rothbart^{1,b)}

¹Center for Epigenetics, Van Andel Research Institute, 333 Bostwick Avenue, NE, Grand Rapids, Michigan 49503, USA

²Laboratory of Structural Sciences, Van Andel Research Institute, 333 Bostwick Avenue, NE, Grand Rapids, Michigan 49503, USA

³VARI/SIMM Center, CAS-Key Laboratory of Receptor Research, Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Shanghai 201203, China

(Received 22 July 2016; accepted 30 September 2016; published online 19 October 2016)

In this communication we introduce an efficient implementation of adaptive biasing that greatly improves the speed of free energy computation in molecular dynamics simulations. We investigated the use of accelerated simulations to inform on compound design using a recently reported and clinically relevant inhibitor of the chromatin regulator BRD4 (bromodomain-containing protein 4). Benchmarking on our local compute cluster, our implementation achieves up to 2.5 times more force calls per day than plumed2. Results of five 1 μ s-long simulations are presented, which reveal a conformational switch in the BRD4 inhibitor between a binding competent and incompetent state. Stabilization of the switch led to a -3 kcal/mol improvement of absolute binding free energy. These studies suggest an unexplored ligand design principle and offer new actionable hypotheses for medicinal chemistry efforts against this druggable epigenetic target class. *Published by AIP Publishing.* [<http://dx.doi.org/10.1063/1.4964776>]

I. INTRODUCTION

Molecular dynamics (MD) simulations are often invoked within the context of drug discovery as a means to score ligands by their absolute binding affinity.¹ In practice, MD simulations are also very useful in the discovery process if they contribute testable hypotheses that guide exploration of the chemical space around a ligand.^{2,3} To be influential, simulations need to provide actionable information on a short time scale, such that a team of medicinal chemists cannot make the same discovery empirically or find an alternative route to success before simulation results are in hand.

The ability to explore the binding process and to compute the free energy of binding within a relevant time frame *in silico* depends on the rate at which the governing equations of motion can be propagated (Figure 1). Standard MD is impractical for describing ligand-protein binding/unbinding because more than 0.5×10^9 force evaluations must be made before the simulation reaches microsecond time scales. Meanwhile, effective ligands may stay bound to their targets for tens of minutes.⁴ With commonly used simulation tools, it takes about a week to simulate a microsecond of a protein-ligand system (Figure 3). Therefore, a considerable gap exists between the rate at which one can perform ligand binding simulations and the rate at which empirically driven progress can be made in a drug discovery setting.

A recent approach to tackle this problem of time scales has been the construction of very specialized computers,^{5,6} which

has generated some long trajectories of ligand binding.^{7,8} These specialized computers aim to make as many MD cycles per unit time as possible. An alternative, more easily dispersed approach is the use of algorithms that leverage formalities of statistical mechanics to extract as much information as possible from every iteration of the MD cycle. This approach aims to reduce the number of cycles required to estimate thermodynamic quantities such as free energy of binding.

In this study, we focus on the use of adaptive biasing potential (ABP) to increase the efficiency of free energy calculations and conformational exploration. ABP methods supplement the system energy in a time-dependent manner so that metastable states are slowly flooded with energy until they are flattened away. The flooding leads to new forces that are contributed by a time evolving biasing potential V_b (Figure 1).

In the MD cycle depicted in Figure 1, ABP methods introduce computational overhead in evaluation of the biasing force, and thus slow the rate of completing the MD cycle. While native MD engines are typically faster at making cycles, ABP schemes improve the statistical value of each cycle by several orders of magnitude. The exact improvement of statistical efficiency can be expressed as the ratio of unbiased to biased partition functions, which quantifies exactly how the ABP schemes improve sampling efficiency by shrinking the Boltzmann-weighted volume of configuration space.

Given the statistical advantage brought about by ABP, the inefficiency of the associated MD cycle is often taken for granted. Here we show that by minimizing ABP network communications within the MD engine and by minimizing computational complexity of the ABP scheme, the efficiency of the biased MD cycle can be drastically improved while

^{a)}Electronic address: bradley.dickson@vai.org

^{b)}scott.rothbart@vai.org




FIG. 1. A basic schematic of the MD propagation cycle before (left) and after (right) the addition of an adaptive biasing potential. The number of trips around the cycle per unit time expresses the speed of MD simulation.

maintaining the statistical advantage of ABP. On our local compute cluster, we show efficiency gains of 25%–67% over the commonly used plumed2⁹ plugin. In Amazon’s elastic cloud computing environment, we demonstrate a 36% improvement in the MD cycle rate for a 26% reduction of simulation costs compared to the same plumed2 plugin. These improvements reduce the time gap between *in silico* and *in vitro* drug discovery, and increase the affordability of high throughput cloud applications. In light of the expanding ABP usage base, these efficiency gains can also be projected to generate substantial savings and an improvement in simulation scope for the computational community.

As an example application, the results of 4 μ s-long simulations of a ligand-protein system are reported for a clinically relevant inhibitor¹⁰ against the chromatin regulator BRD4 (bromodomain-containing protein 4). Three of these simulations were allowed to sample the full surface of the protein and 60% of the simulation volume to fully detail the protein-ligand interaction. These simulations show several metastable states where the ligand aggregates on the protein surface and suggest that better control over a ligand conformation change could improve affinity. The current ligand scaffolds present two conformations, only one of which is binding competent.

Without the efficiency gains derived below, each one of the μ s-long simulations presented in this work would have taken 5.3 days longer to complete on exactly the same computer hardware had the current gold standard ABP software been used. Had all of the simulations been run in serial, they would have taken 21 days extra to complete. In Amazon’s elastic cloud computing environment, the simulations presented below would cost an additional \$1200 USD without the time-saving advances described here. Extrapolating these savings to the MD community projects years of saved compute time and tens of thousands of dollars. Thus, the implementation described below positively impacts computational throughput and better optimizes the dollars spent funding computational research.

II. RESULTS AND DISCUSSION

The efficiency of the ABP implementation has been improved by considering two sources of computational overhead: (1) *computational complexity of the bias potential*. In our implementation, the master node of the simulation

handles the bias updates and bias force evaluation which is a concern for load balancing. Thus, a major advance in this implementation model has been the reduction of computational complexity within the bias itself. We have written the bias in terms of kernels, or “hills,” that have compact support (Figure S1). These new hills result in a significant reduction of operations (namely reduced memory accesses and multiplies) when the bias is updated, leading to better load balancing. Others have also used truncated kernels to reduce complexity¹¹ and below we demonstrate that this one aspect of the implementation is not enough to ensure good scalability. (2) *Minimization of communications introduced by the ABP scheme*. In ABP schemes, a collective variable (CV) must be computed. The CV is a coarse-grained description of the system state, and as such, it requires atomic positions. In most modern MD engines, the master node does not know all of the current atomic positions, so some communication is required. Additionally, once the master node computes the biasing force, that force must be communicated back to the nodes that are managing atoms in the CV. Details regarding each of these aspects of development are given in the [supplementary material](#).

We implemented two bias potentials, with the default being mollified adaptive biasing potential (mABP).¹² This bias potential was derived elsewhere^{12,13} and has been shown to recover well-tempered metadynamics¹⁴ under an approximation that destroys the mollifying properties of the bias.¹³ Importantly, mABP is not a form of metadynamics. Reference 13 shows that while mABP defines the bias potential as the log of a histogram of sampling, all forms of metadynamics write the bias potential as directly proportional to a histogram of sampling. Following from this, a key attribute of mABP is that the biasing force is an ensemble average quantity. The metadynamics family tends toward higher computational complexity than mABP. Boundary corrections,^{15,16} for example, are not required in mABP and are not supported in this distribution. Convergence behavior for mABP and well-tempered metadynamics (WTmetaD) is reported in [supplementary material](#) for alanine dipeptide to test functionality with the compact hills (Figure S2).

It is known that metadynamics places strict conditions on the hill functions¹⁷ and that the convergence time of metadynamics scales exponentially with the hill width.¹³ We observed a parameter dependent sensitivity to artifacts for the well-tempered metadynamics method when using the

hill functions introduced in this work (Figure S3). Artifacts could be observed when WTmetaD was used to repeat the simulations described below as well, though the artifacts were much more difficult to diagnose. Therefore, we do not recommend invoking the metadynamics method with the hills implemented here. WTmetaD performs well for some hill settings but not for others, leaving open the possibility of undiagnosed failures when using the method.

A. Timing of the implementation for a practical application

An application to ligand binding is given here to measure overhead of the new implementation and as a demonstration of the mABP method. The model system involves a recently reported, clinically relevant, small molecule which binds the first bromodomain of BRD4¹⁰ (PDB:5hm0). BRD4 is an acetyllysine binding protein that is emerging as a therapeutic target for controlling expression of the “undruggable” oncogenic transcription factor MYC.¹⁸ Inhibitors like the one simulated here are part of an emerging class of epigenetic-based therapeutics being tested clinically for the treatment of haematological malignancies.¹⁹ The lead ligand (compound 3 of Ref. 10) in complex with bromodomain 1 of BRD4 is shown in Figure 2.

A pair of RMSD CVs were chosen for the simulations. One CV is the RMSD between a subset of atoms in the ligand at its current (simulated) position and the X-ray structure. The second CV is the RMSD between a different subset of ligand atoms in its current (simulated) position and the X-ray. This set of CVs projects the full ligand-accessible space onto the




FIG. 2. (Left) PDB entry 5hm0 is shown with the ligand in green sticks. (Right) The 5hm0 entry is shown again where a clipping plane is used to expose the entire ligand.

diagonal of the CV space because the two CVs are tightly coupled (Figure S4). The base GROMACS input topologies and tpr files for equilibration were built according to the swissparam tutorial for preparing GROMACS topologies.²⁰ Everything required to re-run these biased simulations (all simulation parameters and inputs), including a general tutorial for using the ABP code, can be found at the GitHub page for this project.²¹ The tutorial itself allows one to perform simulations of this ligand-protein system.

The performance of our ABP implementation is compared against native GROMACSv5.0.5 and native GROMACSv5.1 for several numbers of compute cores. Because the implementation builds on the GROMACS package, we call it FABMACS for fast Adaptively Biased Machine for Chemical

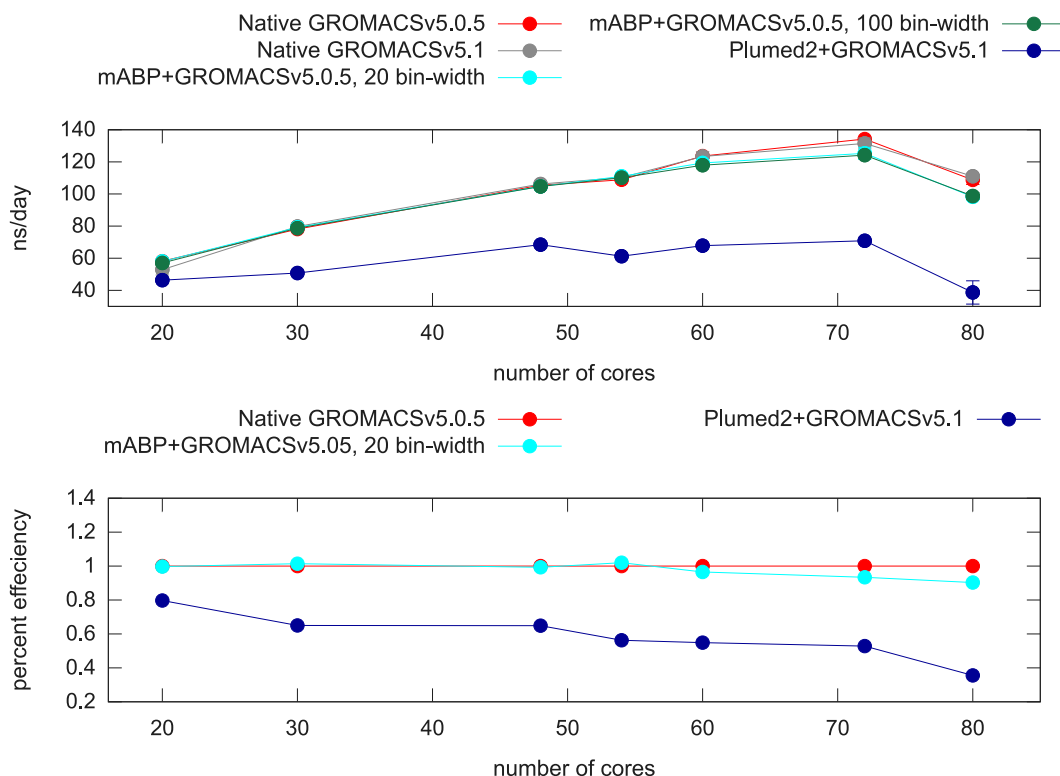


FIG. 3. (Top) Nanoseconds per day as a function of compute cores for GROMACS, mABP (FABMACS), and plumed2. (Bottom) Percent efficiency relative to native GROMACS as a function of cores. All data points are an average of three runs.

Simulations. The data shown in Figure 3 were selected to minimize oscillations in the scaling of the native GROMACS codes, representing the envelope of best performance for the native GROMACS MD engine. In all cases, GROMACS was allowed to choose the partitioning of PME and particle-particle nodes, and in all cases we verified that all the data at a particular node count were collected with the same partitioning. For the fABMACS simulations, the CV space was discretized into a 480 by 480 bin grid with an RMSD range from 0 to 12 nm. Although irrelevant for timing, the biasing parameters were $c = 0.001$ and $b = 0.9$. A range of hill widths were considered, and the temperature was 300 K. Results are reported in units of nanoseconds per day, which can be converted to the number of MD cycles per day by multiplying by 500 000 (the MD timestep was 0.002 fs). See section 1 of the [supplementary material](#) for the definition of the hill functions, S_a^p . The parameter a fixes the width of the hill and p alters the shape. Section III of the [supplementary material](#) covers how to think about these hills in terms of the more common Gaussian hill.

Figure 3 also shows the performance of the plumed2⁹ plugin running WTmetaD, which was plugged into GROMACSv5.1. The fABMACS parameters were converted to well-tempered metadynamics parameters according to Ref. 13 and the hill width used with plumed was within the range of widths shown for fABMACS. The CV and CV discretizations were identical for the plumed2 and fABMACS methods. All of the various GROMACS codes were compiled with the same configurations on the same hardware. We did not observe inefficiency when running WTmetaD in fABMACS. WTmetaD introduces three extra multiplies per grid point during the bias update (to scale hill heights for the potential and its forces), and one exponential evaluation (to determine the hill height). On 72 cores, our implementation of WTmetaD (without boundary corrections) averaged only 1 ns/day less than mABP.

Our ABP implementation presents minimal overhead to the underlying GROMACS engine, and the scalability of GROMACS is intact. For this system, fABMACS achieves a maximum of 125 ns/day on 72 compute cores spanning three different machines. The maximum of the native GROMACS code is 134 ns/day on 72 cores. It can be appreciated, as expected, that the extra ABP-related communications are slowing down the simulation as the number of cores is increased (Figure 3, bottom).

Given the rise of cloud computing, it is pertinent to consider timings in the cloud environment as well. GROMACSv5.1, plumed2, and fABMACS were timed in Amazon’s EC2 environment using “g2.8xlarge” compute nodes. These nodes have 32 cores and four Nvidia GRID GPUs. The cost of this EC2 instance is currently \$2.60 USD per hour. Table I shows the cost per microsecond for native GROMACS, fABMACS, and plumed2 based on timings of 100 000 MD steps. Note that both plumed2 and fABMACS should be expected to generate better sampling than native GROMACS, due to the shrinking partition function, even though native GROMACS is the least expensive per force call. Running fABMACS costs \$38 USD more per microsecond than native GROMACS, while plumed2 costs an extra \$326

TABLE I. Amazon elastic compute cloud results, cost in USD/ μ s, and days/ μ s.

Software	Biased	ns/day	\$(USD)/ μ s	Days/ μ s
GROMACSv5.1	No	82.027	760	12.2
fABMACS	Yes	78.136	798	12.8
Plumed2	Yes	57.449	1086	17.4

USD per μ s. We did not run simulations spanning multiple machines in the cloud because projected costs did not suggest any benefit.

Encouraged by the performance of the application tested above, we also tested the performance of the built-in NAMDv2.11²³ and AMBERv16.05^{24,25} implementations of adaptive biasing potentials for a larger universal benchmark. The benchmark system is a small peptide in water, producing a 99k atom system.²² All collective variables were 2 RMSDs to an initial position with each CV having four atoms. Figure 4 reports results similar to what others have observed for unbiased simulations.^{26,27} Here, we have used 28 cores spread across 7 nodes and to gain insight into the performance of the different MD packages, we watched the network communications on the master node of the simulations via ifconfig. The cumulative inbound traffic is shown in Figure 4 (bottom) for each MD package. All machine interconnects were 40 GbE.

Figure 4 (bottom) suggests that PLUMED2 is inefficient within the bias update and bias force computations but not in network communications. This may be a load balance issue. The built-in NAMD biasing seems limited by the underlying speed of NAMD. The pmemd-AMBER implementation appears to produce excessive network traffic when ABMD is invoked. Activation of biasing within AMBER on GPU also incurs efficiency losses either because excessive system information is passed from the GPU to CPU for biasing at every time step or because CPU speed is a bottleneck. As we noted earlier, this demonstrates that minimization of operational complexity in the bias¹¹ is insufficient to achieve good scalability although it is required for load balancing.

B. Applications to ligand binding

The fABMACS package can be compiled for different scenarios. To keep the code as streamlined as possible, a number of options are hard-coded via a patching script. The script sets a number of parameters for things like CV space discretization, selection of mABP or WTmetaD, and what type of restraint system to apply. There are two selectable restraint types, a spherical or cylindrical restraint. There is a third restraint that cannot be deselected, which limits the maximum value of the CVs. This restraint can be obsoleted by setting a large maximum limit. A full tutorial is available on the project GitHub page,²¹ which walks through setting up and running the simulations presented in this section. The periodic simulation boundaries are unwrapped such that the ligand is whole and such that the ligand does not “jump” across the box. Unwrapping is not supported in some ABP




FIG. 4. Universal benchmark using 99137 atoms for adaptive biasing of two 4-atom CVs, in the canonical ensemble with partial mesh Ewald and periodic boundaries.²² (TOP) Nanoseconds per day for different MD packages on CPU or GPU. (BOTTOM) Network traffic on a node for MD packages when distributed across 7 compute nodes.

plugins,²⁸ yet unwrapping is critical unless the simulation system is designed such that the CVs are never broken by periodicity.

We performed a total of four 1 μ s simulations with the lead ligand that used the spherical (Figure 5(a)) or cylindrical (Figure 5(b)) restraint. The mABP simulations used hills with compact support, called S_a^p , where $p = 20$ and $a = 10 \times \Delta\xi_i$ or $p = 20$ and $a = 20 \times \Delta\xi_i$, where the latter is labeled “wide” in Figure 5(d). Plots of free energy for all of the simulations are shown in 2-dimensions (Figure S4) and as 1-dimensional projections (Figure 5(d)). We followed the formalisms of Refs. 29 and 30 for estimating absolute binding free energy with and without restraints, respectively.

The absolute free energy of binding estimated from each of the simulations ranges from -5.18 to -7.63 kcal/mol. All of the estimates are within $4 k_B T$ of one another, where k_B is Boltzmann’s constant and T is temperature. This is a surprising degree of similarity given that three of the simulations were allowed to sample the entire protein surface, and two of those started in an unbound conformation.

In each of the simulations using wide hills and the large spherical restraint, the crystallographic ligand pose was found only once. One of these simulations started in the bound conformation, one started in an unbound conformation. Both simulations found the crystallographic pose from an unbound state. In the simulation using narrow hills and starting in an unbound conformation, the crystallographic ligand pose was found twice. (Figure S6) Simulations as loosely restrained as these are very rarely performed and have even been described as impossible.²⁹ Indeed, we only tried these simulations after being encouraged by the performance of fABMACS. The simulation using the cylindrical restraint found the crystallographic pose four times. Thus, the free energy estimate from the cylindrical restraint is expected to be most accurate.

The average absolute free energy of binding over the four estimates is -6.56 kcal/mol with a standard error of 0.56 kcal/mol. The free energy estimate from simulation with the cylindrical restraint was -6.16 kcal/mol. We did not compare to binding free energy estimated from IC₅₀ reported

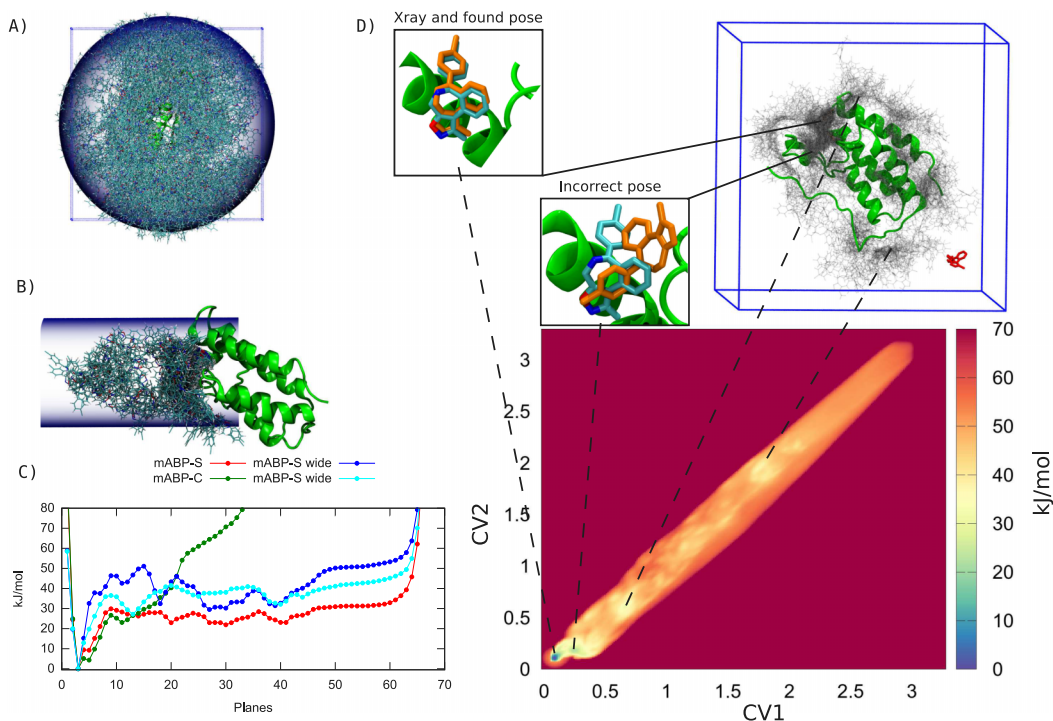


FIG. 5. (a) Snapshots of the ligand from a $1 \mu\text{s}$ simulation with a spherical restraint. The simulation box is shown in blue. (b) Snapshots of the ligand from a $1 \mu\text{s}$ simulation with a cylindrical restraint. (c) 1D projections of the free energy onto CV_1 . The “-S” indicates that the spherical restraint was used while “-C” indicates that the cylindrical restraint was used. The narrow and wide designations are described in the main text. (d) Snapshots of the ligand where at least one ligand atom is within 3.5 \AA of the protein, highlighting some features of the landscape. Simulated ligand poses are shown in orange in the inset panels. Initial ligand position is shown in red.

in Ref. 10 for two reasons. First, we did not parameterize the ligand via *ab initio* methods. Second, the binding constant of the competitor ligand must be known to make an accurate free energy estimate,³¹ and details of the competitor ligand were not reported.

It is also worth noting that most ligand development efforts focus on dissociation constants rather than free energy of binding. If simulation estimates leave a 1 kcal/mol error in the free energy, there is a 5-fold error in the dissociation constant. For a 2 kcal/mol error in the free energy, there is a 29-fold error in the dissociation constant. In general, it is more informative to generate structure-activity

trends than to compare simulation results to experimental metrics.

The simulations using the spherical restraint sample a number of configurations where the ligand is aggregated to the protein surface (Figure 5(c)). We observe a large number of states where the ligand is aggregated on the ZA-loop or docked in the binding pocket with an incorrect pose. Most binding pathways show the ligand aggregate to the pocket or ZA-loop and then explore different ligand poses and azepine ring conformations until the crystallographic pose is found. Often, the ligand leaves the pocket area and moves back into solvent before finding the correct pose.

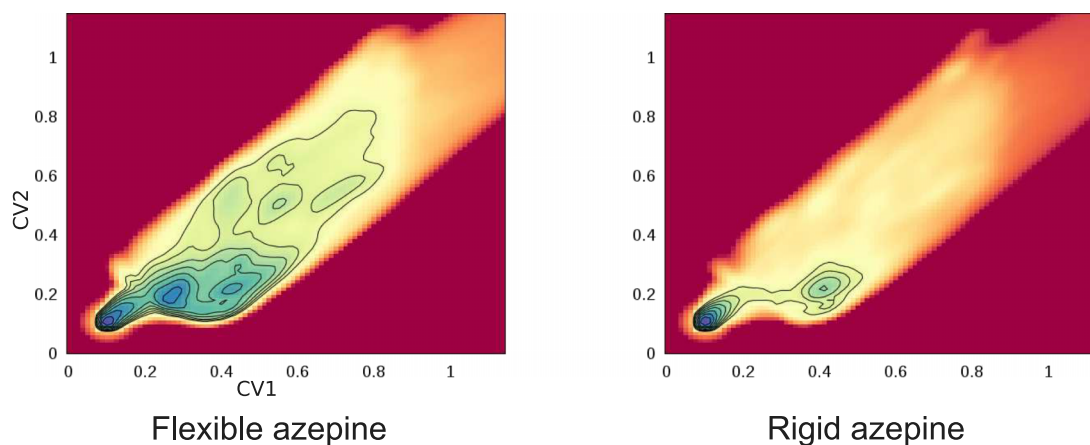


FIG. 6. Binding landscapes without (left) and with (right) a rigid azepine. Contours are every $2k_B T$ from 0 to 40 kJ/mol.

Many attempted binding events were rejected by the protein because the ligand adopted a binding incompetent conformation. The benzoisoxazoloazepine ligand contains an azepine ring that can adopt different puckering states; (Figure S5) only one of the two major puckering states is consistent with the crystallographic pose. The “incorrect pose” shown in Figure 5(c) is representative of a state in which the azepine ring has adopted a conformation inconsistent with the crystallographic pose, precluding correct binding and slowing the binding process.

As a simple means of testing the hypothesis that binding could be improved by reducing flexibility of the azepine ring, we computed the binding free energy again while applying a dihedral restraint that prevents the azepine from changing conformations. The absolute binding free energy, computed from a 1 μ s simulation using the cylindrical restraint and an inflexible azepine, was -9.42 kcal/mol. This compares to -6.16 kcal/mol for the flexible azepine in the cylindrical restraint (Figure 6). The restrained ligand found the crystallographic pose 6 times (Figure S6). These simulations validate our hypothesis and motivate ligand chemistry that would restrain the azepine. Functionalization of the ligand or abandonment of the azepine ring are both plausible directions.

III. CONCLUSION

We have introduced and benchmarked an implementation of ABP in GROMACS 5.0.5, called fABMACS, that presents minimal overhead to the GROMACS engine. The implementation strives to maintain the scalability of GROMACS. A minimal communication scheme and local smoothing kernels have been used to achieve this goal. Drastic efficiency gains were demonstrated, as compared to the most common alternative ABP plugin.

The fABMACS code is available on GitHub²¹ where a tutorial can also be found. All of the ligand-protein simulations presented above can be re-run by following the tutorial. Currently, only RMSD CVs are supported, but a range of systems should be compatible with this. A future goal is to develop the distribution to support a range of different CVs. Cylindrical, spherical, and simple range-limiting CV restraints are supported. A build script is included with the code, which customizes source code for user specifications of methods, restraint types, restraint scales, and so-on. Thus, if a user can install GROMACS 5.0.5 on their compute infrastructure, then fABMACS can be installed using the same configuration. Simply run the build script with specified options and compile the source code as normal. Implementation into other MD engines is also ongoing.

The results of long simulations of the first bromodomain of BRD4 with a clinically relevant small molecule were presented, both to make methodological demonstrations using mABP and to exemplify structural insights that are obtained from atomistic simulations. Simulations were presented that found the bound state after having escaped, even when the full surface of the protein could be sampled. We suggest a trend in the published SAR data that could be

further exploited to improve potency, namely, the azepine ring affords conformational freedom that reduces ligand efficiency.

fABMACS will be useful in the identification of false positives, creation of chemical hypotheses, scoring ligands, and guiding NMR based fragment screening. The efficiency gains of fABMACS will enable more efficient use of computational resources and make new, ambitious applications more affordable and tractable.

SUPPLEMENTARY MATERIAL

See [supplementary material](#) for details on the hill functions, MPI calls in GROMACS, convergence results for alanine dipeptide, free energy landscapes for ligand binding simulations, free energy of azepine conformations, and the definition of our RMSD coordinate.

ACKNOWLEDGMENTS

All computations were performed on the Van Andel Research Institute compute cluster, except for the timings on EC2. This work was supported in part by the Van Andel Research Institute and research grants from the National Institutes of Health to S.B.R. (Grant No. CA181343).

- ¹A. Perez, J. A. Morrone, C. Simmerling, and K. A. Dill, *Curr. Opin. Struct. Biol.* **36**, 25 (2016).
- ²M. A. Camerino, N. Zhong, A. Dong, B. Dickson, L. James, B. Baughman, J. Norris, D. Kireev, W. P. Janzen, and C. Arrowsmith, *Med. Chem. Commun.* **4**, 1501 (2013).
- ³J. I. Stuckey, B. M. Dickson, N. Cheng, Y. Liu, J. L. Norris, S. H. Cholensky, W. Tempel, S. Qin, K. G. Huber, and C. Sagum, *Nat. Chem. Biol.* **12**, 180 (2016).
- ⁴R. A. Copeland, *Nat. Rev. Drug Discovery* **15**, 87 (2015).
- ⁵D. E. Shaw, M. M. Deneroff, R. O. Dror, J. S. Kuskin, and R. H. Larson, *Commun. ACM* **51**, 91 (2008).
- ⁶D. E. Shaw, J. Grossman, J. A. Bank, B. Batson, J. A. Butts, J. C. Chao, M. M. Deneroff, R. O. Dror, A. Even, and C. H. Fenton, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (IEEE Press, 2014), pp. 41–53.
- ⁷Y. Shan, E. T. Kim, M. P. Eastwood, R. O. Dror, M. A. Seeliger, and D. E. Shaw, *J. Am. Chem. Soc.* **133**, 9181 (2011).
- ⁸R. O. Dror, A. C. Pan, D. H. Arlow, D. W. Borhani, P. Maragakis, Y. Shan, H. Xu, and D. E. Shaw, *Proc. Natl. Acad. Sci. U. S. A.* **108**, 13118 (2011).
- ⁹G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, *Comput. Phys. Commun.* **185**, 604 (2014).
- ¹⁰B. K. Albrecht, V. S. Gehling, M. C. Hewitt, R. G. Vaswani, A. Côté, Y. Leblanc, C. G. Nasveschuk, S. Bellon, L. Bergeron, and R. Campbell, *J. Med. Chem.* **59**, 1330 (2016).
- ¹¹V. Babin, C. Roland, and C. Sagui, *J. Chem. Phys.* **128**, 134101 (2008).
- ¹²B. M. Dickson, *Phys. Rev. E* **84**, 037701 (2011).
- ¹³B. M. Dickson, *J. Chem. Phys.* **143**, 234109 (2015).
- ¹⁴A. Barducci, G. Bussi, and M. Parrinello, *Phys. Rev. Lett.* **100**, 020603 (2008).
- ¹⁵M. McGovern and J. de Pablo, *J. Chem. Phys.* **139**, 084102 (2013).
- ¹⁶J. F. Dama, G. M. Hocky, R. Sun, and G. A. Voth, *J. Chem. Theory Comput.* **11**, 5638 (2015).
- ¹⁷J. F. Dama, M. Parrinello, and G. A. Voth, *Phys. Rev. Lett.* **112**, 240602 (2014).
- ¹⁸M. W. Moyer, *Nat. Med.* **17**, 1325 (2011).
- ¹⁹A. Chaidos, V. Caputo, and A. Karadimitris, *Ther. Adv. Hematol.* **6**, 128 (2015).
- ²⁰V. Zoete, M. A. Cuendet, A. Grosdidier, and O. Michielin, *J. Comput. Chem.* **32**, 2359 (2011).
- ²¹B. Dickson, <https://github.com/BradleyDickson/fABMACS>, 2016.
- ²²P. de Waal, <https://github.com/ParkerdeWaal/MD-Universal-Bench>, 2016.

- ²³J. C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kale, and K. Schulten, *J. Comput. Chem.* **26**, 1781 (2005).
- ²⁴R. Salomon-Ferrer, D. A. Case, and R. C. Walker, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **3**, 198 (2013).
- ²⁵R. Salomon-Ferrer, A. W. Gotz, D. Poole, S. Le Grand, and R. C. Walker, *J. Chem. Theory Comput.* **9**, 3878 (2013).
- ²⁶A. H. Poghosyan, G. A. Yeghiazaryan, H. H. Gharabekyan, and A. A. Shahinyan, *Comput. Phys. Commun.* **1**, 736 (2006).
- ²⁷B. Hess, C. Kutzner, D. Van Der Spoel, and E. Lindahl, *J. Chem. Theory Comput.* **4**, 435 (2008).
- ²⁸B. G, <https://groups.google.com/forum/#!msg/plumed-users/Mm9gLNbtznY/CiBbvdyBQAJ>, 2015.
- ²⁹V. Limongelli, M. Bonomi, and M. Parrinello, *Proc. Natl. Acad. Sci. U. S. A.* **110**, 6358 (2013).
- ³⁰I. Buch, T. Giorgino, and G. De Fabritiis, *Proc. Natl. Acad. Sci. U. S. A.* **108**, 10184 (2011).
- ³¹X. Huang, *J. Biomol. Screening* **8**, 34 (2003).