Path space Markov chain Monte Carlo methods for molecular simulation

Thesis by
Jorge L. Rosa-Raíces

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy in Chemistry

Caltech

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2022
Defended April 28, 2022

Jorge L. Rosa-Raíces
ORCID: 0000-0003-2311-2948

Acknowledgments

This thesis is the product of efforts and experiences well beyond academic in scope. Throughout my graduate studies I received copious support and encouragement through multiple avenues, and acknowledgments are due to those responsible.

I am thankful to my advisor, Tom Miller, for his inspiring dedication as a scientist, educator, and leader. During my entire time as a graduate student, Tom provided a perfect balance between stalwart direction, kind encouragement, and truly unwavering support for our common and, most remarkably, my own academic and professional goals. I will forever be indebted to him for his patience and understanding, and for offering within his group a safe space where I could tumble while searching for a sense of belonging in the Caltech community.

For graciously extending beyond his role as collaborator to provide mentorship and guidance in a similar capacity to Tom, I wholeheartedly thank Professor Nawaf Bou-Rabee at Rutgers–Camden. I hope that our future interactions will continue to yield fruitful scientific exploits and foster a long-standing friendship.

During my early years in the Miller group I benefited greatly from the kindness of senior colleagues Matt Welborn, Brett Savoie, Michiel Niesen, Ioan-Bogdan Magdău, and Philip Shushkov, all of whom, intentionally or otherwise, served as role models and offered invaluable career advice I hold dear to this day. Thanks are also due to my cohort of contemporary graduates in the group with whom I was able to grow, comprising Sebastian Lee, Matthew Zimmer, Jeongmin Kim, Xuecheng Tao, Dan Jacobson, Roman Korol, Sherry Cheng, Jiace Sun, Zhuoran Qiao, and Doris Mai; they all provided a pleasant work environment and a forum for interesting discussion of matters scientific and beyond, selflessly stepped in for others in difficult times, and motivated each other to strive further through their impressive accomplishments. It is great to see that our most recent members, namely Jorge Emiliano Deustua, Tomislav Begušić, Rui Li, Vignesh Bhethanabotla, and Bo Li, will carry on the group's spirit by exhibiting these qualities. To all friends and colleagues, mentioned and otherwise, I wish only the best.

Lastly, I am immensely grateful to my family for their love and unconditional support throughout my pursuit of higher education. There is no one more important to me than all of you, and I can only hope to repay any fraction of all you have done for me. Thank you.

Abstract

Path space Markov chain Monte Carlo (McMC) provides a versatile framework for simulating the structure and dynamics of condensed-phase systems aptly described by classical and quantum Boltzmann statistics. This thesis comprises our efforts to design, analyze, and improve path space McMC algorithms to achieve numerically advantageous, and physically accurate, simulation of molecular processes across a range of scales. To improve molecular dynamics (MD) simulations of atomically resolved systems exhibiting pronounced nuclear quantum effects, we introduce a family of integrators for non-preconditioned path-integral MD exhibiting dimension-free statistical accuracy and efficiency, and enabling a many-fold increase in time-step stability relative to conventional approaches at no additional computational cost or implementation complexity. The integrators come with robust performance guarantees that are borne out in thermostatted ring-polymer MD simulations of realistic condensed-phase models. Concurrently, toward extending the range of accessible timescales in stochastic MD simulations of mesoscale coarse-grained molecular systems, we introduce a parallel-in-time integrator for the overdamped Langevin equation based on McMC evaluation of a path-integral representation of the many time-step stochastic MD transition kernel. The parallel-in-time integrator achieves simultaneous integration of multiple stochastic MD time-steps at no greater wall-time cost and with no lesser accuracy than a standard Euler–Maruyama integrator does in serial, and thus instantiates new opportunities to accelerate stochastic dynamics simulations on massively parallel computer architectures. Our work along these two methodological avenues extends the utility of path space McMC across applications in molecular simulation and has broader implications in other disciplines that require accurate and efficient simulations of Markov diffusion processes in state spaces or path spaces.

Published Content and Contributions

[3] J. L. Rosa-Raíces, J. Sun, N. Bou-Rabee, and T. F. Miller III, "A generalized class of strongly stable and dimension-free T-RPMD integrators", The Journal of Chemical Physics **154**, 024106 (2021) `10.1063/5.0036954`.

J.L.R.-R. participated in conceiving the project, participated in the formulation, analysis, and implementation of some of the methods introduced, prepared the simulations to illustrate the methods, and contributed to writing the article.

[2] R. Korol, J. L. Rosa-Raíces, N. Bou-Rabee, and T. F. Miller III, "Dimension-free path-integral molecular dynamics without preconditioning", The Journal of Chemical Physics **152**, 104102 (2020) `10.1063/1.5134810`.

J.L.R.-R. prepared the simulations to illustrate some of the methods introduced, participated in the formulation, analysis, and implementation of some of the methods introduced, and contributed to writing the article.

[1] J. L. Rosa-Raíces, B. Zhang, and T. F. Miller III, "Path-accelerated stochastic molecular dynamics: Parallel-in-time integration using path integrals", The Journal of Chemical Physics **151**, 164120 (2019) `10.1063/1.5125455`.

J.L.R.-R. participated in designing and implementing the methods introduced, prepared the simulations to illustrate the methods, and contributed to writing the article.

Table of Contents

List of Illustrations

List of Tables

Chapter 1

Introduction

This thesis addresses the expansive field of *molecular simulation*, which comprises an ever-growing palette of computational tools for studying the structure and dynamics of molecular systems through the evaluation of principled models. Molecular simulation offers the prospect to observe molecular processes in greater spatiotemporal resolution than is possible with modern experimental techniques, prospectively yielding insights into the structure and dynamics of matter at length-scales ranging from the electronic to the macroscopic [1]; into macromolecular conformational transition rates and their underlying mechanisms [2]; into chemical reactions mediated by complex environments [3]; and into the dynamical response properties of materials in relation to their microstructural fluctuations. Contingent on sufficient modeling accuracy, as increasingly provided by the upscaling of *ab initio* electronic structure techniques *via* data-driven and machine-learning avenues [4, 5], molecular simulation can therefore help predict and interpret experimental outcomes, as well as guide the design of materials and technologies toward the betterment of society.

Despite impressive progress [6–8], it could be argued that the potential of molecular simulation to deliver on the above prospects remains largely unrealized; this is especially so for molecular processes governed by strong and long-ranged quantum-mechanical effects [9, 10] and that involve long time-scale, large length-scale, and mechanistically complex conformational and chemical transformations [11, 12]. The state of affairs calls for further development of theoretical models and computational methods, implementable on suitably optimized hardware [13, 14], that advance the accurate molecular simulation of processes exhibiting these traits. The current thesis aims to facilitate the application of path space Markov chain Monte Carlo (McMC) methods toward this overarching goal.

Besides its modern applications in Bayesian inference [15–17], mathematical finance [18], and various areas of statistical and machine learning [19], path space McMC has proven to be a versatile framework for molecular simulation as demonstrated by the long-standing success of transition path sampling (TPS) [20–23] and path-integral molecular dynamics (PIMD) [24–28]. Like McMC simulation schemes formulated on the state space (e.g., position space or phase space) of a molecular model, such as stochastically thermostatted molecular dynamics (MD) [29], Metropolis–Hastings Monte Carlo (MC) [30], and hybrid MD/MC schemes [31, 32], path space McMC entails the Markovian, step-wise evolution of a *path*, or an ordered collection of replicas of the molecular model, in accordance with a discretized equation of motion or update rule designed to ensure that admissible *path configurations* are visited in accordance with prescribed relative probabilities in the long-time limit [21, 33]. The prescribed stationary distribution of path configurations is determined by the features of the underlying molecular model, as illustrated in Ref. [20] in the case of TPS and in Ref. [28] in the case of PIMD. While recasting a state-space McMC sampling problem as a path-space McMC sampling problem might appear computationally disadvantageous due to the intrinsically higher dimensionality of the latter, it can result in enhanced performance as evidenced by the success of (*i*) non-preconditioned PIMD at estimating the effects of nuclear quantization (as manifest in, e.g., zero-point energy motion and quantum tunneling) on the structural and dynamical properties of condensed-phase systems [3, 10], and of (*ii*) diffusive TPS at generating mechanistically insightful transition path ensembles between metastable conformational states that would be more computationally demanding to discover via direct diffusive MD simulations [22, 34]. Employing molecular paths as variables to be sampled *via* McMC nevertheless entails unique methodological challenges, and the current thesis collects some of our efforts to address them.

Part I of this thesis concerns non-preconditioned PIMD integration, where path-discretization artifacts had previously encumbered the performance and limited the stability of conventional time-discretization schemes [35]. Colleagues recently observed that conventional non-preconditioned PIMD integrators are incompatible with structural features unique to the path space distributions they intend to sample, resulting in a sharp deterioration of their performance and stability under path-discretization refinements [16, 36, 37]. Through the work detailed in Chapters 2 and 3, we introduce a host of *dimension-free* integration schemes for the non-preconditioned PIMD equations of motion that are free of stability-limiting artifacts in both toy-model and realistic applications. The dimension-free integrators are furthermore designed to yield greater long-time statistical accuracy, higher long-time sampling efficiency as measured by the asymptotic variance of observables of practical interest, and faster convergence to the stationary Boltzmann–Gibbs path distribution than conventional schemes. Having found their way into widely available molecular simulation libraries [38, 39], our dimension-free non-preconditioned PIMD integrators are already extending the utility of path space McMC toward elucidating the effects of nuclear quantization in the structure and dynamics of molecular processes modeled at all-atom resolution with *ab initio* potential energy landscapes [40].

Part II of this thesis explores how path space McMC can enhance long time-scale MD simulation of mesoscale coarse-grained molecular systems. Direct sampling of the trajectory space might appear a convenient avenue to evaluate a long MD trajectory at a lower wall-time cost than with direct integration, but a naïve implementation would be hindered by the slow mixing of McMC schemes when sampling long paths [41, 42]. To effectively apply path space McMC to accelerate the integration of long time-scale stochastic MD, Chapter 4 introduces the *path-accelerated molecular dynamics* (PAMD) scheme, which enables parallel-in-time integration of a long stochastic MD trajectory through McMC-based evaluation of a relatively short path distribution governing the near-term evolution of a system over the course of several MD time-steps. Compared to competing parallel-in-time integration schemes for stochastic MD [43, 44], which are based on a prediction-correction paradigm [45] and whose performance can be limited by the quality of the prediction employed, the PAMD scheme shows greater potential for speedup due to its McMC-based formulation, which admits implementations through a plethora of path space sampling methods that can efficiently leverage massively parallel computing architectures. The implementation of PAMD outlined in Chapter 4 is numerically shown to substantially outperform a standard serial-in-time integration scheme for overdamped Langevin dynamics in simple systems, and encourages efforts toward refined implementations amenable to real-world applications in molecular simulation and other disciplines.

References Cited in Chapter 1

[1] R. Radhakrishnan, “A survey of multiscale modeling: Foundations, historical milestones, current status, and future prospects”, AIChE Journal **67**, e17026 (2021).

[2] F. Noé and E. Rosta, “Markov models of molecular kinetics”, The Journal of Chemical Physics **151**, 190401 (2019).

[3] J. S. Kretchmer, N. Boekelheide, J. J. Warren, J. R. Winkler, H. B. Gray, and T. F. Miller III, “Fluctuating hydrogen-bond networks govern anomalous electron transfer kinetics in a blue copper protein”, Proceedings of the National Academy of Sciences **115**, 6129 (2018).

[4] O. A. von Lilienfeld and K. Burke, “Retrospective on a decade of machine learning for chemical discovery”, Nature Communications **11**, 4895 (2020).

[5] H. Kulik, T. Hammerschmidt, J. Schmidt, S. Botti, M. A. L. Marques, M. Boley, M. Scheffler, M. Todorović, P. Rinke, C. Oses, A. Smolyanyuk, S. Curtarolo, A. Tkatchenko, A. Bartok, S. Manzhos, M. Ihara, T. Carrington, J. Behler, O. Isayev, M. Veit, A. Grisafi, J. Nigam, M. Ceriotti, K. T. Schütt, J. Westermayr, M. Gastegger, R. Maurer, B. Kalita, K. Burke, R. Nagai, R. Akashi, O. Sugino, J. Hermann, F. Noé, S. Pilati, C. Draxl, M. Kuban, S. Rigamonti, M. Scheidgen, M. Esters, D. Hicks, C. Toher, P. Balachandran, I. Tamblyn, S. Whitelam, C. Bellinger, and L. M. Ghiringhelli, “Roadmap on machine learning in electronic structure”, Electronic Structure, Accepted Manuscript (2022).

[6] R. O. Dror, R. M. Dirks, J. P. Grossman, H. Xu, and D. E. Shaw, “Biomolecular simulation: A computational microscope for molecular biology”, Annual Review of Biophysics **41**, 429 (2012).

[7] N. E. Jackson, M. A. Webb, and J. J. de Pablo, “Recent advances in machine learning towards multiscale soft materials design”, Current Opinion in Chemical Engineering **23**, 106 (2019).

[8] F. Noé, A. Tkatchenko, K.-R. Müller, and C. Clementi, “Machine learning for molecular simulation”, Annual Review of Physical Chemistry **71**, 361 (2020).

[9] M. Ceriotti, W. Fang, P. G. Kusalik, R. H. McKenzie, A. Michaelides, M. A. Morales, and T. E. Markland, “Nuclear quantum effects in water and aqueous systems: Experiment, theory, and current challenges”, Chemical Reviews **116**, 7529 (2016).

[10] T. E. Markland and M. Ceriotti, “Nuclear quantum effects enter the mainstream”, Nature Reviews Chemistry **2**, 0109 (2018).

[11] R. Elber, “Perspective: Computer simulations of long time dynamics”, The Journal of Chemical Physics **144**, 060901 (2016).

[12] D. E. Makarov, “Life in silico: Are we close yet?”, Proceedings of the National Academy of Sciences **118**, e2100278118 (2021).

[13] D. E. Shaw, R. O. Dror, J. K. Salmon, J. P. Grossman, K. M. Mackenzie, J. A. Bank, C. Young, M. M. Deneroff, B. Batson, K. J. Bowers, E. Chow, M. P. Eastwood, D. J. Ierardi, J. L. Klepeis, J. S. Kuskin, R. H. Larson, K. Lindorff-Larsen, P. Maragakis, M. A. Moraes, S. Piana, Y. Shan, and B. Towles, “Millisecond-scale molecular dynamics simulations on Anton”, in Proceedings of the conference on high performance computing networking, storage and analysis, SC ’09 (2009), pp. 1–11.

[14] J. P. Grossman, B. Towles, B. Greskamp, and D. E. Shaw, “Filtering, reductions and synchronization in the Anton 2 network”, Parallel and Distributed Processing Symposium, 860 (2015).

[15] A. M. Stuart, “Inverse problems: A Bayesian perspective”, Acta Numerica **19**, 451 (2010).

[16] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White, “MCMC methods for functions: Modifying old algorithms to make them faster”, Statistical Science **28**, 424 (2013).

[17] M. Dashti and A. M. Stuart, “The Bayesian Approach to Inverse Problems”, in *Handbook of uncertainty quantification*, edited by R. Ghanem, D. Higdon, and H. Owhadi (Springer International Publishing, 2017), pp. 311–428.

[18] V. Linetsky, “The path integral approach to financial modeling and options pricing”, Computational Economics **11**, 129 (1997).

[19] G. W. Benton, W. J. Maddox, J. P. Salkey, J. Albinati, and A. G. Wilson, “Function-space distributions over kernels”, in *Proceedings of the 33rd international conference on neural information processing systems* (Curran Associates, Inc., Red Hook, NY, 2019), pp. 1–12.

[20] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, “Transition path sampling: Throwing ropes over rough mountain passes, in the dark”, Annual Reviews of Physical Chemistry **53**, 291 (2002).

[21] M. G. Reznikoff and E. Vanden-Eijnden, “Invariant measures of stochastic partial differential equations and conditioned diffusions”, Comptes Rendus Mathematique **340**, 305 (2005).

[22] T. F. Miller III and C. Predescu, “Sampling diffusive transition paths”, The Journal of Chemical Physics **126**, 144102 (2007).

[23] F. J. Pinski and A. M. Stuart, “Transition paths in molecules at finite temperature”, The Journal of Chemical Physics **132**, 184104 (2010).

[24] D. Chandler and P. G. Wolynes, “Exploiting the isomorphism between quantum theory and classical statistical mechanics of polyatomic fluids”, The Journal of Chemical Physics **74**, 4078 (1981).

[25] M. Parrinello and A. Rahman, “Study of an F center in molten KCl”, The Journal of Chemical Physics **80**, 860 (1984).

[26] I. R. Craig and D. E. Manolopoulos, “Quantum statistics and classical mechanics: Real time correlation functions from ring polymer molecular dynamics”, The Journal of Chemical Physics **121**, 3368 (2004).

[27] T. F. Miller III and D. E. Manolopoulos, “Quantum diffusion in liquid water from ring polymer molecular dynamics”, The Journal of Chemical Physics **123**, 154504 (2005).

[28] S. Habershon, D. E. Manolopoulos, T. E. Markland, and T. F. Miller III, “Ring-polymer molecular dynamics: Quantum effects in chemical dynamics from classical trajectories in an extended phase space”, Annual Reviews of Physical Chemistry **64**, 387 (2013).

[29] N. Bou-Rabee, “Time integrators for molecular dynamics”, Entropy **16**, 138 (2014).

[30] A. Sokal, “Monte Carlo methods in statistical mechanics: Foundations and new algorithms”, in *Functional integration: Basics and applications*, edited by C. DeWitt-Morette, P. Cartier, and A. Folacci (Springer Science & Business Media, 1997), pp. 131–192.

[31] J. P. Nilmeier, G. E. Crooks, D. D. L. Minh, and J. D. Chodera, “Nonequilibrium candidate Monte Carlo is an efficient tool for equilibrium simulation”, Proceedings of the National Academy of Sciences **108**, E1009 (2011).

[32] N. Bou-Rabee and J. M. Sanz-Serna, “Geometric integrators and the Hamiltonian Monte Carlo method”, Acta Numerica **27**, 113 (2018).

[33] M. Hairer, A. M. Stuart, and J. Voss, “Analysis of SPDEs arising in path sampling part II: The nonlinear case”, The Annals of Applied Probability **17**, 1657 (2007).

[34] J. Lee, I.-H. Lee, I. Joung, L. Jooyoung, and B. R. Brooks, “Finding multiple reaction pathways via global optimization of action”, Nature Communications **8**, 15443 (2017).

[35] R. Korol, N. Bou-Rabee, and T. F. Miller III, “Cayley modification for strongly stable path-integral and ring-polymer molecular dynamics”, The Journal of Chemical Physics **151**, 124103 (2019).

[36] A. Beskos, G. Roberts, A. Stuart, and J. Voss, “MCMC methods for diffusion bridges”, Stochastics and Dynamics **8**, 319 (2008).

[37] N. Bou-Rabee, “Cayley splitting for second-order Langevin stochastic partial differential equations”, Preprint (2017).

[38] A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in 't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, and S. J. Plimpton, "LAMMPS: a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales", Computer Physics Communications **271**, 108171 (2022).

[39] V. Kapil, M. Rossi, O. Marsalek, R. Petraglia, Y. Litman, T. Spura, B. Cheng, A. Cuzzocrea, R. H. Meißner, D. M. Wilkins, B. A. Helfrecht, P. Juda, S. P. Bienvenue, W. Fang, J. Kessler, I. Poltavsky, S. Vandenbrande, J. Wieme, C. Corminboeuf, T. D. Kühne, D. E. Manolopoulos, T. E. Markland, J. O. Richardson, A. Tkatchenko, G. A. Tribello, V. Van Speybroeck, and M. Ceriotti, "i-PI 2.0: A universal force engine for advanced molecular simulations", Computer Physics Communications **236**, 214 (2019).

[40] D. L. Eldridge, R. Korol, M. K. Lloyd, A. C. Turner, M. A. Webb, T. F. Miller III, and D. A. Stolper, "Comparison of experimental vs theoretical abundances of $^{13}CH_3D$ and $^{12}CH_2D_2$ for isotopically equilibrated systems from 1 to 500 °C", ACS Earth and Space Chemistry **3**, 2747 (2019).

[41] M. Hairer, A. M. Stuart, J. Voss, and P. Wiberg, "Analysis of SPDEs arising in path sampling. Part I: The Gaussian case", Communications in Mathematical Sciences **3**, 587 (2005).

[42] T. R. Gingrich and P. L. Geissler, "Preserving correlations between trajectories for efficient path sampling", The Journal of Chemical Physics **142**, 234104 (2015).

[43] A. L. Blumers, Z. Li, and G. E. Karniadakis, "Supervised parallel-in-time algorithm for long-time Lagrangian simulations of stochastic dynamics: Application to hydrodynamics", Journal of Computational Physics **393**, 214 (2019).

[44] F. Legoll, T. Lelièvre, and U. Sharma, "An adaptive Parareal algorithm: Application to the simulation of molecular dynamics trajectories", SIAM Journal on Scientific Computing **44**, B146 (2022).

[45] J.-L. Lions, Y. Maday, and G. Turinici, "A parareal discretization in time of PDE's", Comptes Rendus de l'Académie des Sciences - Series I - Mathematics **332**, 661 (2001).

Part I

DIMENSION-FREE RING-POLYMER MOLECULAR DYNAMICS

Chapter 2

Thermostatted ring-polymer molecular dynamics integrators with dimensionality freedom

Adapted from

R. Korol, J. L. Rosa-Raíces, N. Bou-Rabee, and T. F. Miller III, “Dimension-free path-integral molecular dynamics without preconditioning”, The Journal of Chemical Physics **152**, 104102 (2020) DOI: `10.1063/1.5134810`.

Abstract

Convergence with respect to imaginary-time discretization (i.e., the number of ring-polymer beads) is an essential part of any path-integral-based molecular dynamics (MD) calculation. However, an unfortunate property of existing non-preconditioned numerical integration schemes for path-integral molecular dynamics (PIMD), including essentially all existing ring-polymer molecular dynamics (RPMD) and thermostatted RPMD (T-RPMD) methods, is that for a given MD time-step, the overlap between the exact ring-polymer Boltzmann–Gibbs distribution and that sampled using MD becomes zero in the infinite-bead limit. This has clear implications for “Metropolized” MD sampling schemes such as hybrid Monte Carlo, and it also causes the divergence with bead number of the primitive path-integral kinetic energy expectation value when using standard RPMD or T-RPMD. We show that these and other problems can be avoided through the introduction of “dimension-free” numerical integration schemes for which the sampled ring-polymer position distribution has non-zero overlap with the exact distribution in the infinite-bead limit for the case of a harmonic potential. Most notably, we introduce the BCOCB integration scheme, which achieves dimension freedom via a particular symmetric splitting of the integration time-step and a novel implementation of the Cayley modification [*J. Chem. Phys.* **151**, 124103 (2019)] for the free ring-polymer half-steps. More generally, we show that dimension freedom can be achieved via mollification of the forces from the external physical potential. The dimension-free path-integral numerical integration schemes introduced here yield finite error bounds for a given MD time-step,

even as the number of beads is taken to infinity; these conclusions are proven for the case of a harmonic potential and borne out numerically for anharmonic systems that include liquid water. The numerical results for BCOCB are particularly striking, allowing for nearly three-fold increases in the stable time-step for liquid water with respect to the Bussi-Parrinello (OBABO) and Leimkuhler (BAOAB) integrators while introducing negligible errors in the calculated statistical properties and absorption spectrum. Importantly, the dimension-free, non-preconditioned integration schemes introduced here preserve strong stability, symplecticity, time reversibility, and global second-order accuracy; and they remain simple, black-box methods that avoid additional computational costs, tunable parameters, or system-specific implementations.

2.1 Introduction

Considerable effort has been dedicated to the development of numerical integration schemes for imaginary-time path-integral molecular dynamics (PIMD) [1]. In comparison to standard classical molecular dynamics, PIMD numerical integration faces the additional challenge of the highly oscillatory dynamics of the ring-polymer internal modes. Work on numerical integrators for PIMD generally falls into two distinct categories. In the first, the PIMD equations of motion are *preconditioned* by modifying the ring-polymer mass matrix [2–10]; this approach, which includes the widely used staging algorithms [11], causes the integrated trajectories to differ from those of the ring-polymer molecular dynamics (RPMD) model for real-time dynamics [12, 13], but it can lead to efficient [4–6] sampling of the quantum Boltzmann–Gibbs distribution [14, 15]. In the second category, no modification is made to the ring-polymer mass matrix, i.e., the equations of motion are *non-preconditioned* [13, 16–21].

With the aim of providing useful models for real-time quantum dynamics, as well as simple and efficient algorithms for equilibrium thermal sampling, the current chapter focuses on non-preconditioned PIMD numerical integration, notable examples of which include RPMD [12, 13] and its thermostatted variant T-RPMD [20]. Numerical integration schemes for these methods typically employ symmetric factorizations of the time-evolution operator of the form [11, 16–25]

$$\begin{aligned} \mathrm{e}^{\Delta t\mathcal{L}_n} \approx \mathrm{e}^{a\frac{\Delta t}{2}\mathcal{O}_n}\mathrm{e}^{\frac{\Delta t}{2}\mathcal{B}_n}\mathrm{e}^{\frac{\Delta t}{2}\mathcal{A}_n}\mathrm{e}^{(1-a)\Delta t\mathcal{O}_n}\mathrm{e}^{\frac{\Delta t}{2}\mathcal{A}_n}\mathrm{e}^{\frac{\Delta t}{2}\mathcal{B}_n}\mathrm{e}^{a\frac{\Delta t}{2}\mathcal{O}_n} \\ \text{with} \;\; a \in \{0,1\} \; , \end{aligned} \tag{2.1}$$

where the operator $\mathcal{L}_n = \mathcal{A}_n + \mathcal{B}_n + \mathcal{O}_n$ includes contributions from the purely harmonic free ring-polymer motion ($\mathcal{A}_n$), the external potential gradient ($\mathcal{B}_n$), and a thermostat ($\mathcal{O}_n$), and where n is the (typically large) number of beads or modes comprising the ring polymer. Note that the standard microcanonical RPMD numerical integration scheme is recovered in the limit of zero coupling to the thermostat, and that Eq. (2.1) yields the "OBABO" scheme of Bussi and Parrinello [22] when $a = 1$ and the "BAOAB" scheme of Leimkuhler [25] when $a = 0$.

In a recent publication [26], it was emphasized that earlier PIMD numerical integration schemes had overlooked a fundamental aspect of the $\exp((\Delta t/2)\mathcal{A}_n)$ sub-step of the time evolution in Eq. (2.1). Standard practice in these integration schemes has been to exactly evolve the harmonic free ring-polymer dynamics associated with $\exp((\Delta t/2)\mathcal{A}_n)$ using the uncoupled free ring-polymer normal modes [11, 16–18], which was shown to lack the property of strong stability in the numerical integration, leading to resonance instabilities for microcanonical RPMD and loss of ergodicity for T-RPMD [26]. Use of the Cayley modification to the free ring-polymer motion was shown to impart strong stability to the time-evolution, thereby improving numerical stability for microcanonical RPMD and restoring ergodicity for T-RPMD [26].

In the current study, we focus on the accuracy of both statistical and dynamical properties of the OBABO and BAOAB schemes, as well as the corresponding integrators obtained when the exact free ring-polymer step is replaced by the strongly stable Cayley modification (OBCBO and BCOCB, respectively). Particular attention is paid to the effect of finite-time-step error with these integrators in the limit of large bead numbers. Of these four integrators, it is found that only BCOCB is *dimension-free*, in the sense that the sampled ring-polymer position distribution has non-zero overlap with the exact distribution in the infinite-bead limit for the case of a harmonic potential. Further, it is demonstrated that the OBCBO scheme can be made dimension-free via the technique of force mollification. It is shown that the newly introduced BCOCB integrator yields better accuracy than all other considered non-preconditioned PIMD integrators and allows for substantially larger time-steps in the calculation of both statistical and dynamical properties. Importantly, these gains are made without loss of computational efficiency or algorithmic simplicity.

2.2 Non-preconditioned PIMD

Consider a one-dimensional quantum particle with the Hamiltonian operator

$$\hat{H} \;=\; \frac{1}{2m}\,\hat{p}^2 + V(\hat{q}) \;, \tag{2.2}$$

where m is the particle mass, $\hat{q}$ and $\hat{p}$ the position and momentum operators, and $V(\hat{q})$ a potential energy surface. Ignoring exchange statistics, the properties of this system at thermal equilibrium are encoded in the quantum partition function

$$Q \;=\; \mathrm{Tr}[\mathrm{e}^{-\beta\hat{H}}] \;, \tag{2.3}$$

where $\beta = (k_B T)^{-1}$, k_B is the Boltzmann constant and T the thermodynamic temperature. Using a path-integral discretization (i.e., a Trotter factorization of the Boltzmann–Gibbs operator [27]), $Q \;=\; \lim_{n\to\infty} Q_n$ can be approximated by the classical partition function Q_n of a ring polymer with n beads [14, 15],

$$Q_n \;=\; \frac{m^n}{(2\pi\hbar)^n}\int \mathrm{d}^n\boldsymbol{q}\int \mathrm{d}^n\boldsymbol{v}\;\mathrm{e}^{-\beta H_n(\boldsymbol{q},\boldsymbol{v})} \;, \tag{2.4}$$

where $\boldsymbol{q} = [q_0,\ldots,q_{n-1}]^{\mathrm{T}}$ is the vector of bead positions and $\boldsymbol{v}$ the corresponding vector of velocities. The ring-polymer Hamiltonian is given by

$$H_n(\boldsymbol{q},\boldsymbol{v}) \;=\; H_n^0(\boldsymbol{q},\boldsymbol{v}) + V_n^{\mathrm{ext}}(\boldsymbol{q}) \;, \tag{2.5}$$

which includes contributions from the physical potential

$$V_n^{\mathrm{ext}}(\boldsymbol{q}) \;=\; \frac{1}{n}\sum_{j=0}^{n-1} V(q_j) \tag{2.6}$$

and the free ring-polymer Hamiltonian

$$H_n^0(\boldsymbol{q},\boldsymbol{v}) \;=\; \frac{m_n}{2}\sum_{j=0}^{n-1}\left[v_j^2 + \omega_n^2(q_{j+1}-q_j)^2\right] \;, \tag{2.7}$$

where $m_n = m/n$, $\omega_n = n/(\hbar\beta)$ and $q_n = q_0$.

Non-preconditioned PIMD evolves the phase $[\boldsymbol{q}(t)^{\mathrm{T}},\boldsymbol{v}(t)^{\mathrm{T}}]^{\mathrm{T}}$ of the ring polymer as per

$$\begin{aligned}\dot{\boldsymbol{q}}(t) \;=\; \boldsymbol{v}(t)\;;\;\dot{\boldsymbol{v}}(t) \;&=\; -\,\boldsymbol{\Omega}^2\boldsymbol{q}(t) + m_n^{-1}\boldsymbol{F}(\boldsymbol{q}(t)) \\ &\quad -\boldsymbol{\Gamma}\boldsymbol{v}(t) + \sqrt{2\beta^{-1}m_n^{-1}}\boldsymbol{\Gamma}^{1/2}\dot{\boldsymbol{W}}(t) \;,\end{aligned} \tag{2.8}$$

which corresponds to the Hamiltonian dynamics of $H_n(\boldsymbol{q}, \boldsymbol{v})$ with an Ornstein–Uhlenbeck thermostat. In Eq. (3.1) we introduced $\boldsymbol{F}(\boldsymbol{q}) = -\nabla V_n^{\text{ext}}(\boldsymbol{q})$, an n-dimensional standard Brownian motion $\boldsymbol{W}(t)$ and the $n \times n$ matrices

$$\begin{aligned} \boldsymbol{\Omega} &= \boldsymbol{U} \operatorname{diag}(0, \omega_{1,n}, \dots, \omega_{n-1,n}) \boldsymbol{U}^{\mathrm{T}} \text{ and} \\ \boldsymbol{\Gamma} &= \boldsymbol{U} \operatorname{diag}(0, \gamma_1, \dots, \gamma_{n-1}) \boldsymbol{U}^{\mathrm{T}} , \end{aligned} \tag{2.9}$$

where $\gamma_j \geq 0$ is the jth friction coefficient, $\boldsymbol{U}$ an $n \times n$ real orthogonal discrete Fourier transform matrix such that

$$\boldsymbol{\Omega}^2 = -\omega_n^2 \begin{bmatrix} -2 & 1 & 0 & \cdots & 0 & 1 \\ 1 & -2 & 1 & 0 & \cdots & 0 \\ & \ddots & \ddots & \ddots & & \\ & & \ddots & \ddots & \ddots & \\ 0 & \cdots & 0 & 1 & -2 & 1 \\ 1 & 0 & \cdots & 0 & 1 & -2 \end{bmatrix} , \tag{2.10}$$

and the $\omega_{j,n}$, with $j = 0, \dots, n-1$, are the so-called Matsubara frequencies [28]

$$\omega_{j,n} = \begin{cases} 2\omega_n \sin\left(\frac{\pi j}{2n}\right) & \text{if } j \text{ is even} , \\ 2\omega_n \sin\left(\frac{\pi(j+1)}{2n}\right) & \text{else} . \end{cases} \tag{2.11}$$

Observe that the zero-frequency (i.e., centroid) ring-polymer mode is uncoupled from the thermostat, and the coefficients $\{\gamma_j\}_{j=1}^{n-1}$ in Eq. (2.9) constitute the friction schedule applied to the non-centroid modes.

The standard method for discretizing Eq. (2.8) is to use a symmetric splitting method of the form of Eq. (2.1) that consists of a combination of three types of sub-steps: (*i*) exact free ring-polymer evolution of time-step τ,

$$\begin{bmatrix} \boldsymbol{q} \\ \boldsymbol{v} \end{bmatrix} \leftarrow \exp(\tau \boldsymbol{A}) \begin{bmatrix} \boldsymbol{q} \\ \boldsymbol{v} \end{bmatrix} \quad \text{with} \quad \boldsymbol{A} = \begin{bmatrix} \boldsymbol{0} & \boldsymbol{I} \\ -\boldsymbol{\Omega}^2 & \boldsymbol{0} \end{bmatrix} , \tag{2.12}$$

where $\boldsymbol{A}$ is the Hamiltonian matrix associated to the free ring polymer, (*ii*) velocity updates of time-step τ due to forces from the external potential,

$$\boldsymbol{v} \leftarrow \boldsymbol{v} + \tau m_n^{-1} \boldsymbol{F}(\boldsymbol{q}) , \tag{2.13}$$

and (*iii*) velocity updates of time-step τ due to the thermostat,

$$\boldsymbol{v} \leftarrow \exp(-\tau \boldsymbol{\Gamma}) \boldsymbol{v} + \sqrt{\beta^{-1} m_n^{-1}} (\boldsymbol{I} - \exp(-2\tau \boldsymbol{\Gamma}))^{1/2} \boldsymbol{\xi} , \tag{2.14}$$

where $\boldsymbol{I}$ is the $n \times n$ identity matrix and $\boldsymbol{\xi}$ is an n-dimensional vector whose components are independent, standard normal random variables. The acronyms "OBABO" and "BAOAB" indicate the order in which these sub-steps are applied, as specified in Eq. (2.1) with $a = 1$ or $a = 0$, respectively.

In previous work [26], some of us showed that the matrix exponential for the free ring polymer evolution in Eq. (2.12) is not a strongly stable symplectic matrix, and that as a consequence, the OBABO and BAOAB schemes can display non-ergodicity at time-steps $\Delta t = k\pi/\omega_{j,n}$ for any $1 \leq j \leq n$ and $k \geq 1$. We also identified a maximum safe time-step size $\Delta t_{\star} = \beta\hbar\pi/(2n)$, below which the matrix exponential is strongly stable. As $n \to \infty$, this maximum safe time-step goes to zero, such that no finite time-step for the scheme in Eq. (2.1) is safe in this limit from non-ergodicity.

This non-ergodicity motivates the Cayley modification [26] which consists of approximating the matrix exponential appearing in Eq. (2.12) with the Cayley transform. Specifically, for the Cayley-modified OBABO scheme (called OBCBO), we replace the exact free ring-polymer update of time-step $\tau = \Delta t$ with

$$\operatorname{cay}(\Delta t \boldsymbol{A}) = (\boldsymbol{I} - (1/2)\Delta t \boldsymbol{A})^{-1}(\boldsymbol{I} + (1/2)\Delta t \boldsymbol{A}) \ . \tag{2.15}$$

For the Cayley-modified BAOAB scheme (called BCOCB), we replace the two exact free ring-polymer updates of half-time-step $\tau = \Delta t/2$ with $\operatorname{cay}(\Delta t \boldsymbol{A})^{1/2}$. While it might be expected that these half-time-step updates would instead be replaced with $\operatorname{cay}((\Delta t/2)\boldsymbol{A})$, such a choice leads to a loss of strong stability. Our use of the square root of the Cayley transform preserves strong stability, symplecticity, time reversibility, local third-order accuracy, and guarantees that $\operatorname{cay}(\Delta t \boldsymbol{A})^{1/2} \operatorname{cay}(\Delta t \boldsymbol{A})^{1/2} = \operatorname{cay}(\Delta t \boldsymbol{A})$. Furthermore, the square root of the Cayley transform is no more complicated to evaluate than the Cayley transform itself. Both the OBCBO and BCOCB Cayley modifications of Eq. (2.1) are ergodic for a fixed time-step, irrespective of the number of beads; moreover, like Eq. (2.1), the Cayley modified integrators exhibit locally third-order accuracy in the time-step and leave invariant the free ring-polymer Boltzmann–Gibbs distribution in the special case of a constant external potential ($V \equiv$ const.) [26].

2.3 BCOCB avoids pathologies in the infinite-bead limit

In this section, we show that of the OBABO, BAOAB, OBCBO and BCOCB integration schemes, only BCOCB is dimension-free. Although the current section presents analytical results for the specific case of a harmonic external potential, these results are supported by numerical results for anharmonic external potentials in Secs. 2.6 and 2.7.

To this end, consider the jth internal ring-polymer mode with frequency $\omega_{j,n}$, in the presence of a harmonic external potential $V(q) = (1/2)\Lambda q^2$ and a Langevin thermostat with friction γ_j. Expressed in terms of the normal mode coordinates, obtained from the Cartesian positions and velocities via the orthogonal transformation

$$\boldsymbol{\varrho} = \boldsymbol{U}^{\mathrm{T}}\boldsymbol{q} \quad \text{and} \quad \boldsymbol{\varphi} = \boldsymbol{U}^{\mathrm{T}}\boldsymbol{v} \tag{2.16}$$

where $\boldsymbol{U}$ is defined in Eq. (3.2), the non-preconditioned PIMD equations of motion for this mode are

$$\begin{aligned} \begin{bmatrix} \dot{\varrho}_j(t) \\ \dot{\varphi}_j(t) \end{bmatrix} &= \boldsymbol{K}_{j,n} \begin{bmatrix} \varrho_j(t) \\ \varphi_j(t) \end{bmatrix} + \begin{bmatrix} 0 \\ \sqrt{2\beta^{-1} m_n^{-1} \gamma_j} \dot{W}_j(t) \end{bmatrix} \\ \boldsymbol{K}_{j,n} &= \boldsymbol{A}_{j,n} + \boldsymbol{B} + \boldsymbol{O}_j \ , \end{aligned} \tag{2.17}$$

where $\dot{W}_j$ is a scalar white-noise and we have introduced the 2×2 matrices

$$\boldsymbol{A}_{j,n} = \begin{bmatrix} 0 & 1 \\ -\omega_{j,n}^2 & 0 \end{bmatrix} \ , \ \boldsymbol{B} = \begin{bmatrix} 0 & 0 \\ -\Lambda/m & 0 \end{bmatrix} \ \text{and} \ \boldsymbol{O}_j = \begin{bmatrix} 0 & 0 \\ 0 & -\gamma_j \end{bmatrix} \ .$$

The solution $(\varrho_j(t), \varphi_j(t))$ of Eq. (2.17) is a bivariate Gaussian, and in the limit as $t \to \infty$, the probability distribution of $(\varrho_j(t), \varphi_j(t))$ converges to a centered bivariate normal distribution with covariance matrix

$$\boldsymbol{\Sigma}_{j,n} = \frac{1}{\beta m_n} \begin{bmatrix} s_{j,n}^2 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{where} \quad s_{j,n}^2 = \frac{1}{\Lambda/m + \omega_{j,n}^2} \ . \tag{2.18}$$

For this system, a single time-step of Eq. (2.1) can be compactly written as

$$\begin{bmatrix} \varrho_j(t+\Delta t) \\ \varphi_j(t+\Delta t) \end{bmatrix} = \boldsymbol{M}_{j,n,\Delta t} \begin{bmatrix} \varrho_j(t) \\ \varphi_j(t) \end{bmatrix} + \boldsymbol{R}_{j,n,\Delta t}^{1/2} \begin{bmatrix} \xi_0 \\ \eta_0 \end{bmatrix} \ , \tag{2.19}$$

where ξ_0 and η_0 are independent standard normal random variables and we have introduced the 2×2 matrices

$$\begin{aligned} \boldsymbol{M}_{j,n,\Delta t} &= \mathrm{e}^{a\frac{\Delta t}{2}\boldsymbol{O}_j} \mathrm{e}^{\frac{\Delta t}{2}\boldsymbol{B}} \mathrm{e}^{\frac{\Delta t}{2}\boldsymbol{A}_{j,n}} \mathrm{e}^{(1-a)\Delta t \boldsymbol{O}_j} \mathrm{e}^{\frac{\Delta t}{2}\boldsymbol{A}_{j,n}} \mathrm{e}^{\frac{\Delta t}{2}\boldsymbol{B}} \mathrm{e}^{a\frac{\Delta t}{2}\boldsymbol{O}_j} \\ \boldsymbol{R}_{j,n,\Delta t} &= \frac{1 - \mathrm{e}^{-2(1-a)\gamma_j \Delta t}}{\beta m_n} \boldsymbol{N}_{j,n,\Delta t} \boldsymbol{P} \boldsymbol{N}_{j,n,\Delta t}^{\mathrm{T}} \\ &\quad + \frac{1 - \mathrm{e}^{-a\gamma_j \Delta t}}{\beta m_n} \big((\boldsymbol{M}_{j,n,\Delta t} \mathrm{e}^{-a\frac{\Delta t}{2}\boldsymbol{O}_j}) \boldsymbol{P} (\boldsymbol{M}_{j,n,\Delta t} \mathrm{e}^{-a\frac{\Delta t}{2}\boldsymbol{O}_j})^{\mathrm{T}} + \boldsymbol{P} \big) \end{aligned}$$

where $\boldsymbol{P} = \left[\begin{smallmatrix} 0 & 0 \\ 0 & 1 \end{smallmatrix}\right]$ and $\boldsymbol{N}_{j,n,\Delta t} = \mathrm{e}^{a\frac{\Delta t}{2}\boldsymbol{O}_j}\mathrm{e}^{\frac{\Delta t}{2}\boldsymbol{B}}\mathrm{e}^{\frac{\Delta t}{2}\boldsymbol{A}_{j,n,\Delta t}}$. The corresponding step for the Cayley modification is obtained by replacing $\exp((\Delta t/2)\boldsymbol{A}_{j,n})$ in Eq. (2.19) with $\mathrm{cay}(\Delta t\boldsymbol{A}_{j,n})^{1/2}$, which is given by

$$\mathrm{cay}(\Delta t\boldsymbol{A}_{j,n})^{1/2} = \sqrt{\frac{1}{4+\omega_{j,n}^2\Delta t^2}} \begin{bmatrix} 2 & \Delta t \\ -\omega_{j,n}^2\Delta t & 2 \end{bmatrix} . \tag{2.20}$$

A sufficient condition[1] for ergodicity of Eq. (2.19) is

$$1 \ > \ \mathsf{A}_{j,n,\Delta t}^2 \cosh^2((\Delta t/2)\gamma_j) \ , \tag{2.21}$$

where

$$\mathsf{A}_{j,n,\Delta t} \ = \ \cos(\Delta t\omega_{j,n}) - \frac{(\Lambda/m)\Delta t}{2\omega_{j,n}} \sin(\Delta t\omega_{j,n}) \ .$$

For the Cayley modification of Eq. (2.19), Eq. (2.21) still provides a sufficient condition for ergodicity, except with

$$\mathsf{A}_{j,n,\Delta t} \ = \ \frac{4-(\omega_{j,n}^2+2(\Lambda/m))\Delta t^2}{4+\omega_{j,n}^2\Delta t^2} \ .$$

Due to the lack of strong stability in the exact free ring-polymer evolution, Eq. (2.19) fails to meet the condition in Eq. (2.21) and becomes non-ergodic whenever $\Delta t = k\pi/\omega_{j,n}$ where $k \geq 1$ [26]; no such problem exists for the Cayley modification. Regardless, assuming that the condition in Eq. (2.21) holds, the numerical stationary distribution is a centered Gaussian with 2×2 covariance matrix $\boldsymbol{\Sigma}_{j,n,\Delta t}$ that satisfies the linear equation

$$\boldsymbol{\Sigma}_{j,n,\Delta t} = \boldsymbol{M}_{j,n,\Delta t}\boldsymbol{\Sigma}_{j,n,\Delta t}\boldsymbol{M}_{j,n,\Delta t}^{\mathrm{T}} + \boldsymbol{R}_{j,n,\Delta t} \ ,$$

for which the solution is

$$\boldsymbol{\Sigma}_{j,n,\Delta t} = \frac{1}{\beta m_n} \begin{bmatrix} s_{j,n,\Delta t}^2 & 0 \\ 0 & r_{j,n,\Delta t}^2 \end{bmatrix} . \tag{2.22}$$

The variances in the position and velocity marginals of the numerical stationary distribution with the covariance matrix in Eq. (2.22) are, respectively, $(\beta m_n)^{-1}s_{j,n,\Delta t}^2$

---

[1] In the special case when $\Lambda = 0$, the given condition for OBCBO corrects a sign error in Eq. (37) of Ref. [26].

and $(\beta m_n)^{-1} r^2_{j,n,\Delta t}$ with

$$s^2_{j,n,\Delta t} = \begin{cases} \dfrac{1}{\omega^2_{j,n} + \frac{\Lambda \Delta t \omega_{j,n}}{m} \cot(\Delta t \omega_{j,n}) - (\frac{\Lambda \Delta t}{2m})^2} & a = 1 \\ \dfrac{1}{\omega^2_{j,n} + \frac{\Lambda \Delta t \omega_{j,n}}{2m} \cot(\frac{\Delta t}{2} \omega_{j,n})} & a = 0 \end{cases} \tag{2.23}$$

$$r^2_{j,n,\Delta t} = \begin{cases} 1 & a = 1 \\ \dfrac{2m\omega_{j,n} - \Lambda \Delta t \tan(\frac{\Delta t}{2} \omega_{j,n})}{2m\omega_{j,n}} & a = 0 \end{cases} \tag{2.24}$$

For the Cayley modification of Eq. (2.19),

$$s^2_{j,n,\Delta t} = \frac{4m}{4m - a\Delta t^2 \Lambda} s^2_{j,n} \,, \tag{2.25}$$

$$r^2_{j,n,\Delta t} = \frac{4m - (1-a)\Delta t^2 \Lambda}{4m} \,. \tag{2.26}$$

Note that these numerical stationary distributions are independent of the friction parameter γ_j, which is a benefit of schemes based on splitting the non-preconditioned PIMD dynamics into Hamiltonian and thermostat parts, and using the exact Ornstein–Uhlenbeck flow in Eq. (2.14) to evolve the thermostat part. Moreover, comparing the exact covariance matrix in Eq. (2.18) with the finite-time-step approximations in Eqs. (2.22)–(2.26), note that in all cases $\mathbf{\Sigma}_{j,n} = \lim_{\Delta t \to 0} \mathbf{\Sigma}_{j,n,\Delta t}$. These results have previously been reported for the OBABO (Eqs. (2.23) and (2.24), $a = 1$) and BAOAB (Eqs. (2.23) and (2.24), $a = 0$) schemes [8, 29] but not for the OBCBO (Eqs. (2.25) and (2.26), $a = 1$) or BCOCB (Eqs. (2.25) and (2.26), $a = 0$) schemes.

In normal mode coordinates, the exact and numerical position-marginals can be written as an infinite product of one-dimensional centered normal distributions with variances given by $(\beta m_n)^{-1} s^2_{j,n}$ and $(\beta m_n)^{-1} s^2_{j,n,\Delta t}$, respectively. By Kakutani's theorem [30, 31], these two distributions have a non-zero overlap in the infinite-bead limit if and only if the limit series

$$\lim_{n \to \infty} \sum_{j=1}^{n-1} \left(1 - \frac{s_{j,n}}{s_{j,n,\Delta t}} \right)^2 \tag{2.27}$$

converges. For OBABO and BAOAB, due to the oscillatory cotangent term appearing in $s_{j,n,\Delta t}$, the limit $\lim_{n \to \infty} (1 - s_{n-1,n}/s_{n-1,\Delta t})^2$ does not exist, and therefore, the series does not converge. For OBCBO, the jth summand of this series is

$$\frac{\Delta t^4 \Lambda^2}{16 m^2} \left(1 + \sqrt{\frac{4m - \Delta t^2 \Lambda}{4m}} \right)^{-2} \,,$$

which more obviously diverges. Therefore, for OBABO, OBCBO, and BAOAB, the numerical stationary distribution has no overlap with the exact stationary distribution in the infinite-bead limit; it is in this sense that these schemes fail to exhibit the property of dimensionality freedom. Remarkably, BCOCB is exact in the position-marginal and thus exhibits dimensionality freedom. See Sec. 2.9 for a brief summary of the properties of other symmetric splittings that were considered.

2.4 Consequences for the primitive kinetic energy expectation value

In the current section, we show that the non-overlap pathology of the OBABO, BAOAB, and OBCBO schemes causes a divergence with increasing bead number of the primitive path-integral kinetic-energy expectation value, an issue that is numerically well known for OBABO and BAOAB [8, 29, 32, 33]. We further show that this divergence is fully eliminated via the BCOCB scheme, as expected.

The primitive kinetic energy expectation value is given by [34, 35]

$$\begin{aligned}\langle \mathrm{KE}_n^{\mathrm{pri}} \rangle &= \frac{n}{2\beta} - \sum_{j=1}^{n} \frac{m_n \omega_n^2}{2} \langle (q_j - q_{j-1})^2 \rangle && (2.28)\\ &= \frac{1}{2\beta} + \sum_{j=1}^{n-1} \left(\frac{1}{2\beta} - \frac{m_n \omega_{j,n}^2}{2} \langle \varrho_j^2 \rangle \right) && (2.29)\end{aligned}$$

where the first equality involves a sum over the ring-polymer beads in Cartesian coordinates (with $q_n = q_0$) and the second equality performs the summation in terms of the ring-polymer normal modes. The divergence of this expectation value is numerically illustrated for the simple case of a harmonic oscillator (Figs. 2.1a-d); note that for larger MD time-steps, the OBABO, BAOAB, and OBCBO schemes fail to reach a plateau with increasing bead number and dramatically deviate from the exact result (dashed line). The same divergence for OBABO and BAOAB has been numerically observed in many systems [8, 29, 32, 33], including liquid water which we discuss later. A striking observation from Figs. 2.1a-d is that the BCOCB exhibits no such divergence or error in the primitive kinetic energy expectation value at high bead number, regardless of the employed time-step.

Using Eq. (2.18), note that the contribution to the primitive kinetic energy expectation value from the jth ring-polymer mode is

$$\langle \mathrm{KE}_{j,n}^{\mathrm{pri}} \rangle = \frac{1}{2\beta} \left(1 - \omega_{j,n}^2 s_{j,n}^2 \right) \ ,$$

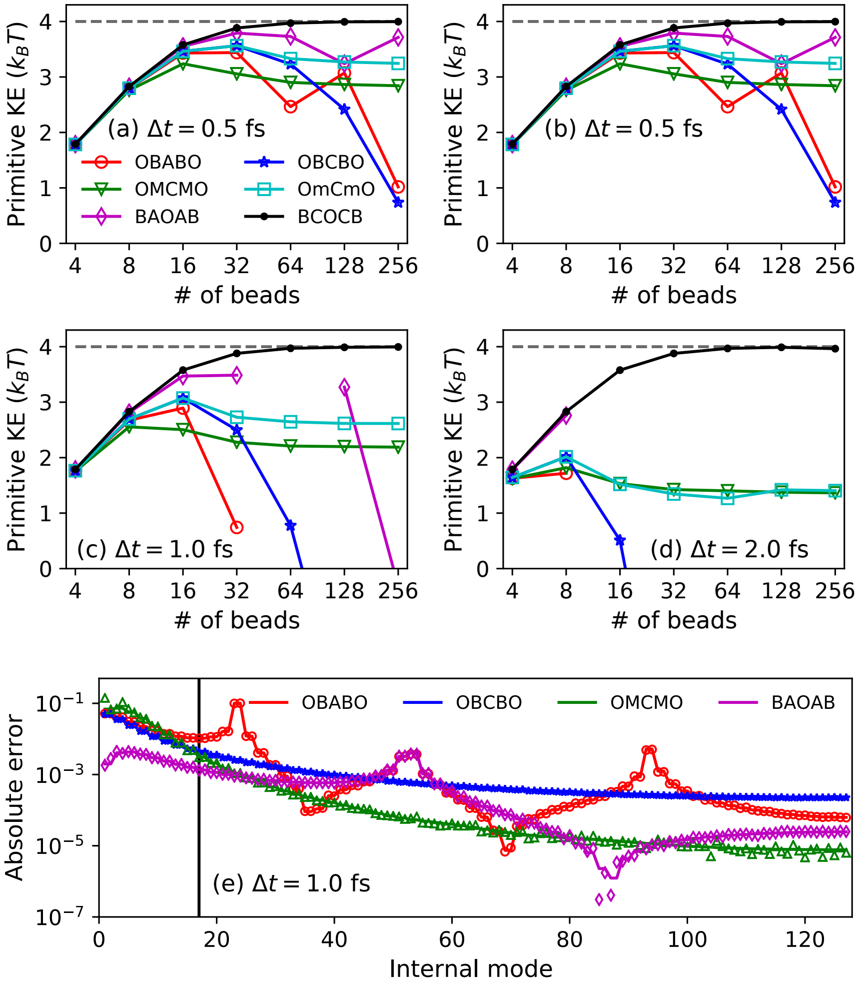

Figure 2.1. Primitive kinetic energy expectation values for a harmonic potential $V(q) = \frac{1}{2}\Lambda q^2$ with force constant $\Lambda = 256$, $\hbar = 1$, mass $m = 1$ and reciprocal temperature $\beta = 1$; choosing energies to be in units of $k_{\mathrm{B}}T$ at room temperature (300 K), $\beta\hbar \approx 25.5$ fs and $\Lambda = m\omega^2$ where $\omega = 3315\,\mathrm{cm}^{-1}$. (Panels a–d) For various MD time-steps Δt, the primitive kinetic energy expectation value as a function of the number of ring-polymer beads, with the exact kinetic energy indicated as a dashed gray line. The standard error of all visible data points in each plot is smaller than the symbol size. (Panel e) Per-mode error in the variance of position coordinate of the normal modes for simulations run with 128 ring-polymer beads and a time-step of 1 fs; solid lines are analytic predictions from Eq. (2.33) with (2.23) and (2.25) defining $s^2_{j,n,\Delta t}$ for the different schemes; points indicate the results of numerical PIMD simulations using the various integration schemes. The BCOCB scheme is not shown since it has zero error for all internal modes. The black vertical line indicates the crossover frequency ($\omega_{\mathrm{x}} = 2/\Delta t$) for the error of OBCBO and OMCMO based on the bounds in Eqs. (2.43) and (2.44).

such that in the infinite-bead limit,

$$\lim_{n\to\infty}\sum_{j=1}^{n-1}\langle \mathrm{KE}^{\mathrm{pri}}_{j,n}\rangle = \frac{\hbar}{4}\sqrt{\frac{\Lambda}{m}}\coth\left(\frac{\beta\hbar}{2}\sqrt{\frac{\Lambda}{m}}\right) \ . \tag{2.30}$$

Similarly using Eq. (2.22), the jth-mode contribution to the kinetic energy from the finite-time-step numerical expectation value is

$$\langle \mathrm{KE}^{\mathrm{pri}}_{j,n,\Delta t}\rangle = \frac{1}{2\beta}\left(1-\omega^2_{j,n}s^2_{j,n,\Delta t}\right) \ . \tag{2.31}$$

Thus, the per-mode error in kinetic energy is

$$|\,\langle \mathrm{KE}^{\mathrm{pri}}_{j,n}\rangle - \langle \mathrm{KE}^{\mathrm{pri}}_{j,n,\Delta t}\rangle\,| = \frac{m_n\omega^2_{j,n}}{2}\rho_{j,n,\Delta t} \ , \tag{2.32}$$

where the per-mode error in the position-marginal for internal mode j is

$$\rho_{j,n,\Delta t} = \frac{1}{\beta m_n}\left|s^2_{j,n} - s^2_{j,n,\Delta t}\right| \ , \tag{2.33}$$

where $s_{j,n,\Delta t}$ is given by Eq. (2.23) for the cases of OBABO ($a = 1$) and BAOAB ($a = 0$) and by Eq. (2.25) for the cases of OBCBO ($a = 1$) and BCOCB ($a = 0$). Note that this error vanishes only for the BCOCB scheme, which satisfies $\rho_{j,n,\Delta t} = 0$ for each mode j, irrespective of the time-step Δt.

Equations (2.32) and (2.33) indicate that the primitive kinetic energy estimator is a sensitive measure of the finite-time-step error in the sampled ring-polymer position distribution associated with the high-frequency modes. Figure 2.1e resolves this per-mode error, $\rho_{j,n,\Delta t}$, for each internal mode in simulations that employ a total of 128 beads, including results from OBABO (red), BAOAB (magenta), and OBCBO (blue), using a time-step of 1 fs, with the solid lines indicating the analytical predictions in Eq. (2.33) and with the dots indicating the result of numerical simulations. The analytical results are fully reproduced by the simulations. Note that the OBABO per-mode error exhibits dramatic spikes for $\omega_{j,n}\Delta t = k\pi$ where $1 \le j \le n$ and for some $k \ge 1$, which coincide with the loss of ergodicity of that integration scheme. The BAOAB scheme exhibits these resonance instabilities at even values of k. However, it is the failure of this per-mode error to sufficiently decay as a function of the mode number for all three of OBABO, BAOAB, and OBCBO that gives rise upon summation to the divergence of the primitive kinetic energy expectation value, as seen for this particular time-step value in Fig. 2.1d. Since $\omega^2_{j,n}s^2_{j,n} \to 1$ as $n \to \infty$, the convergence of $\lim_{n\to\infty}\sum_{j=1}^{n-1}|\,\langle \mathrm{KE}^{\mathrm{pri}}_{j,n}\rangle - \langle \mathrm{KE}^{\mathrm{pri}}_{j,n,\Delta t}\rangle\,|$ reduces to the convergence of the series $\lim_{n\to\infty}\sum_{j=1}^{n-1}\left|s^2_{j,n} - s^2_{j,n,\Delta t}\right|$, which diverges for both OBABO and OBCBO due to the same reasons as discussed in Sec. 2.3.

2.5 Dimensionality freedom for OBCBO via force mollification

Sections 2.3 and 2.4 have demonstrated that whereas the BCOCB integrator exhibits dimensionality freedom, the OBCBO integrator does not. In the current section, we show that this shortcoming of OBCBO can be addressed by the use of force mollification, in which the external potential energy in Eq. (2.6) is replaced by

$$\tilde{V}_n^{\mathrm{ext}}(\boldsymbol{q}) = V_n^{\mathrm{ext}}(\mathrm{sinc}(\tilde{\boldsymbol{\Omega}}\Delta t/2)\boldsymbol{q}) \ , \tag{2.34}$$

where $\tilde{\boldsymbol{\Omega}}$ is any positive semi-definite $n \times n$ matrix that has the same eigenvectors as $\boldsymbol{\Omega}$ (Eq. (3.2)) while possibly having different eigenvalues. Force mollification has not previously been employed for PIMD, although the strategy originates from a variation-of-constants formulation of the solution to Eq. (2.8) [36–39]; specifically, the protocol in Eq. (2.34) is a generalization of the mollified impulse method [36].

The use of force mollification for enhancing non-preconditioned PIMD integration can be motivated on physical grounds: In the absence of a physical potential, four of the considered integration schemes (OBABO, BAOAB, OBCBO, and BCOCB) leave invariant the exact free ring-polymer Boltzmann–Gibbs distribution [26]. Therefore, the loss of any overlap between the exact stationary distribution of the position marginals in the infinite-bead limit for OBABO, BAOAB, and OBCBO must be attributed to the influence of the time-evolution from the external potential in the schemes (i.e., the “B” sub-step) as implemented in Eq. (2.13); the BCOCB scheme does not suffer from this problem. To remove this pathology in the OBCBO scheme, we thus use mollification to “taper down” the external forces on the high-frequency modes, such that the resulting integration correctly reverts to free-particle motion for those modes, which should become decoupled from the external potential as the frequency increases. The specific appearance of the $1/2$ factor in the sinc function argument ensures that the sinc function switches from its high-frequency effect to its low-frequency effect when the period of the Matsubara frequency is commensurate with Δt; the zero-frequency ring-polymer centroid mode is untouched by mollification.

Force mollification requires only a small algorithmic modification of the OBCBO integrator. Specifically, the “B” sub-step in Eq. (2.13) is replaced with

$$\boldsymbol{v} \leftarrow \boldsymbol{v} + \frac{\Delta t}{2} m_n^{-1} \tilde{\boldsymbol{F}}(\boldsymbol{q}) \ , \tag{2.35}$$

where the mollified forces are

$$\tilde{\boldsymbol{F}}(\boldsymbol{q}) = \mathrm{sinc}(\tilde{\boldsymbol{\Omega}}\Delta t/2)\boldsymbol{F}(\tilde{\boldsymbol{q}}) = \boldsymbol{U}\boldsymbol{D}_{\Delta t}\boldsymbol{U}^{\mathrm{T}}\boldsymbol{F}(\tilde{\boldsymbol{q}}) \tag{2.36}$$

where $\tilde{\boldsymbol{q}} = \boldsymbol{U}\boldsymbol{D}_{\Delta t}\boldsymbol{U}^{\mathrm{T}}\boldsymbol{q}$ are the mollified bead positions and where $\boldsymbol{D}_{\Delta t}$ is the diagonal matrix of eigenvalues associated with $\mathrm{sinc}(\tilde{\boldsymbol{\Omega}}\Delta t/2)$, i.e.,

$$\boldsymbol{D}_{\Delta t} = \mathrm{diag}\left(\mathrm{sinc}(\tilde{\omega}_{0,n}\Delta t/2), \ldots, \mathrm{sinc}(\tilde{\omega}_{n-1,n}\Delta t/2)\right) \tag{2.37}$$

where $\tilde{\omega}_{j,n}$ is the jth eigenvalue of $\tilde{\boldsymbol{\Omega}}$. In practice, the mollified forces are computed in normal mode coordinates as follows:

(a) Starting with the ring-polymer bead position in normal mode coordinates, obtain a copy of the mollified bead positions via

$$\tilde{\boldsymbol{q}} = \boldsymbol{U}\boldsymbol{D}_{\Delta t}\boldsymbol{\varrho} \, . \tag{2.38}$$

(b) Evaluate the external forces at the mollified ring-polymer bead positions, $\boldsymbol{F}(\tilde{\boldsymbol{q}})$.

(c) Apply the remaining mollification to the forces in Eq. (2.36) via

$$\boldsymbol{U}^{\mathrm{T}}\tilde{\boldsymbol{F}}(\boldsymbol{q}) = \boldsymbol{D}_{\Delta t}\boldsymbol{U}^{\mathrm{T}}\boldsymbol{F}(\tilde{\boldsymbol{q}}) \, . \tag{2.39}$$

We emphasize that in comparison to the standard force update (Eq. (2.13)) the use of the mollified force update (Eq. (2.35)) introduces neither additional evaluations of the external forces nor $n \times n$ matrix multiplies associated with the discrete Fourier transform; it therefore avoids any significant additional computational cost.

This mollification scheme preserves reversibility and symplecticity as well as local third-order accuracy of the OBCBO scheme with time-step. We emphasize that the sinc-function-based mollification scheme in Eq. (2.35) is not unique and alternatives can certainly be devised. Even within the functional form of the mollification in Eq. (2.35), flexibility remains with regard to the choice of the matrix $\tilde{\boldsymbol{\Omega}}$, which allows for mode-specificity in the way the mollification is applied. A simple choice for this matrix is $\tilde{\boldsymbol{\Omega}} = \boldsymbol{\Omega}$, such that mollification is applied to all of the non-zero ring-polymer internal modes. With this choice, we arrive at a fully-specified integration scheme that replaces the original “B” sub-step in Eq. (2.13) with the mollified-force sub-step in Eq. (2.35); we shall refer to this force-mollified version of OBCBO integration scheme as “OMCMO.” In Sec. 2.5.1, we propose a partially mollified choice for $\tilde{\boldsymbol{\Omega}}$ that further improves accuracy.

For the harmonic external potential, the previously derived relations for OBCBO (most notably Eqs. (2.21), (2.25)–(2.26) and (2.32)–(2.33)) also hold for OMCMO with Λ suitably replaced by $\tilde{\Lambda}_j = \mathrm{sinc}^2(\omega_{j,n}\Delta t/2)\Lambda$. Note that $\tilde{\Lambda}_j \leq \Lambda$, since

$\mathrm{sinc}^2(\mathsf{x}) \leq 1$ for all $\mathsf{x} \geq 0$, making clear that the mollification reduces the effect of the external potential on the higher-frequency internal ring-polymer modes.

We now show that mollifying the forces in the B sub-step fixes the pathologies of OBCBO in the infinite-bead limit, by restoring overlap between the sampled and exact stationary distributions. To see this, note that the jth summand in Eq. (2.27) for OMCMO satisfies

$$\left(1 - \frac{s_{j,n}}{s_{j,n,\Delta t}}\right)^2 \leq \left(1 - \frac{s_{j,n}^2}{s_{j,n,\Delta t}^2}\right)^2 \leq f(\omega_j \Delta t/2)\frac{\Delta t^4 \Lambda^2}{16m^2} \ ,$$

where $f(\mathsf{x}) = \left((1 - \mathrm{sinc}^2(\mathsf{x}))/\mathsf{x}^2 + \mathrm{sinc}^2(\mathsf{x})\right)^2$ and we have used the infinite-bead limit for the ring-polymer internal-mode frequencies

$$\omega_j = \lim_{n\to\infty} \omega_{j,n} = \begin{cases} \dfrac{\pi j}{\hbar\beta} & \text{if } j \text{ is even} \ , \\ \dfrac{\pi (j+1)}{\hbar\beta} & \text{else} \ . \end{cases} \tag{2.40}$$

since[2]

$$\sum_{j=1}^{\infty} f(\omega_j \Delta t/2) \ \leq \ 6\frac{\hbar\beta}{\pi\Delta t} + 4 \ ,$$

we obtain

$$\lim_{n\to\infty} \sum_{j=1}^{n-1} \left(1 - \frac{s_j}{s_{j,n,\Delta t}}\right)^2 \leq \left(6\frac{\hbar\beta}{\pi\Delta t} + 4\right)\frac{\Delta t^4 \Lambda^2}{16m^2} \ . \tag{2.41}$$

Again invoking Kakutani's theorem (Eq. (2.27)), it follows that the numerical stationary distribution has an overlap with the exact stationary distribution. As a byproduct of this analysis, we can also quantify the amount of overlap between the exact

---

[2] This inequality comes from using Eq. (2.40) to write $\sum_{j=1}^{\infty} f(\omega_j \Delta t/2) = \mathrm{I} + \mathrm{II}$ where

$$\mathrm{I} = 2 \sum_{j=1}^{\lfloor \hbar\beta/(\pi\Delta t) \rfloor} f(j\pi\Delta t/(\hbar\beta)) \ \text{ and } \ \mathrm{II} = 2 \sum_{j=\lceil \hbar\beta/(\pi\Delta t) \rceil}^{\infty} f(j\pi\Delta t/(\hbar\beta)) \ .$$

Then the first term admits the bound

$$\mathrm{I} \leq 2f(1)\hbar\beta/(\pi\Delta t) < 4\hbar\beta/(\pi\Delta t) \ ;$$

for the second term, we use

$$\mathrm{II} \leq F(1) + \hbar\beta/(\pi\Delta t) \int_1^{\infty} F(x)\mathrm{d}x$$

where $F(\mathsf{x}) = 2((1 - \mathrm{sinc}^2(\mathsf{x}))/\mathsf{x}^2 + 1/\mathsf{x}^2)^2$ is monotone decreasing on $[1, \infty)$ with $F(1) \leq 4$ and $\int_1^{\infty} F(x)\mathrm{d}x \leq 2$.

and numerically sampled stationary distributions, revealing that the total variation distance [40] between these distributions is given by[3]

$$d_{\mathrm{TV}}(\mu, \mu_{\Delta t}) \le \sqrt{\left(6\frac{\hbar\beta}{\pi\Delta t} + 4\right)\frac{\Delta t^2 \Lambda}{2m}} \,. \tag{2.42}$$

In summary, the force mollification strategy introduced here provably removes the pathologies due to the "B" sub-step in the case of a harmonic oscillator potential. Moreover, for any finite number of beads, the total variation distance between the exact and numerically sampled stationary distribution can be bounded by Eq. (2.42), and thus OMCMO admits dimension-free error bounds.

Before proceeding, we first return to Fig. 2.1 to compare the accuracy of OMCMO with the un-mollified OBCBO scheme for the internal mode position marginal of the harmonic oscillator. As seen in Fig. 2.1e for the results with a time-step of 1 fs, the per-mode error obtained by the mollified scheme (OMCMO, green) decays more rapidly with mode number than does OBCBO. Figure 2.1d further illustrates that upon summation of the per-mode contributions, the OMCMO prediction for the primitive kinetic energy converges to a well-defined asymptote with respect to the number of ring-polymer beads, whereas OBCBO diverges as discussed earlier. Similar behavior is seen for shorter MD time-steps (panels a-c), although the failure of OBCBO becomes less severe with this range of bead numbers as the time-step is reduced.

Although it is satisfying that mollification via OMCMO both formally and numerically ameliorates the problems of the OBCBO scheme in the high-bead-number limit, the OMCMO results in Fig. 2.1 are not ideal, since in some cases the OMCMO error is substantially larger than that of OBCBO when a modest number of beads is used (e.g., for 16 beads in panel d). This observation points to a simple and general refinement of the OMCMO scheme, which we discuss in the next subsection.

---

[3] This quantification uses: (*i*) $d_{\mathrm{TV}} \le 2^{1/2} d_{\mathrm{H}}$ where d_{TV} is the total variation distance and d_{H} is the Hellinger distance; and (*ii*) subadditivity of the squared Hellinger distance, which implies that

$$\begin{aligned} d_{\mathrm{H}}^2(\mu, \mu_{\Delta t}) &\le \lim_{n\to\infty} \sum_{j=1}^{n-1} d_{\mathrm{H}}^2(\mathcal{N}(0, s_{j,n}^2), \mathcal{N}(0, s_{j,n,\Delta t}^2)) \\ &\le \lim_{n\to\infty} \sum_{j=1}^{n-1} (1 - s_{j,n}^2/s_{j,n,\Delta t}^2)^2 \le (3\hbar\beta/(\pi\Delta t) + 2)\,\frac{\Delta t^4 \Lambda^2}{8m^2} \,. \end{aligned}$$

2.5.1 Partial mollification

Comparison of the per-mode errors from OBCBO and OMCMO in Fig. 2.1e reveals that lower errors for OMCMO are only enjoyed for internal modes that exceed a particular frequency (indicated by the vertical black line). This observation suggests that if a "crossover frequency" could be appropriately defined, then a refinement to OMCMO could be introduced for which mollification is applied only to the ring-polymer internal modes with frequencies that exceed this crossover value.

For the case of a harmonic external potential, this crossover frequency ω_{x} can be found by comparing a bound for the per-mode error (Eq. (2.33)) for OBCBO

$$\rho_{j,n,\Delta t} \leq \left(\frac{1}{m_n \omega_{j,n}^2 \beta} \frac{\Delta t^2 \Lambda}{4m - \Delta t^2 \Lambda} \right) \tag{2.43}$$

to that for OMCMO

$$\rho_{j,n,\Delta t} \leq g(\omega_{j,n}\Delta t/2) \left(\frac{1}{m_n \omega_{j,n}^2 \beta} \frac{\Delta t^2 \Lambda}{4m - \Delta t^2 \Lambda} \right) \ , \tag{2.44}$$

where $g(\mathrm{x}) = (1 - \mathrm{sinc}^2(\mathrm{x}))/\mathrm{x}^2 + \mathrm{sinc}^2(\mathrm{x})$. Since $g(\mathrm{x}) \geq 1$ only when $\mathrm{x} \leq 1$, we expect better accuracy if mollification is only applied to those ring-polymer internal modes with frequencies $\omega_{j,n} \geq \omega_{\mathrm{x}}$, where $\omega_{\mathrm{x}} = 2/\Delta t$. Although this result was derived for the case of a harmonic potential, it does not depend on Λ. We call this resulting partly mollified integration scheme "OmCmO." This scheme has the nice properties of OMCMO, including strong stability and dimensionality freedom.

Implementation of OmCmO is a trivial modification of OMCMO, requiring only that the diagonal elements of $\boldsymbol{D}_{\Delta t}$ in Eq. (2.37) are evaluated using

$$\mathrm{sinc}(\tilde{\omega}_{j,n}\Delta t/2) = \begin{cases} 1 & \text{for } \omega_{j,n} < \omega_{\mathrm{x}} \\ \mathrm{sinc}(\omega_{j,n}\Delta t/2) & \text{otherwise} \ , \end{cases} \tag{2.45}$$

where $j = 0, \ldots, n-1$. In physical terms, the emergence of $2/\Delta t$ in the crossover frequency is intuitive, since as was previously mentioned, it corresponds to having the ring-polymer mode undergo a full period per time-step Δt.

Finally, numerical results for the case of a harmonic potential (Figs. 2.1a-d) reveal that the partially modified OmCmO scheme (cyan) achieves both robust convergence of the primitive kinetic energy with increasing bead number, as well as consistently better accuracy than the OBCBO and OMCMO integration schemes, as expected. However, it must be emphasized that for all panels of Fig. 2.1, the BCOCB scheme (which requires no force mollification) is by far the most accurate and stable.

2.6 Results for anharmonic oscillator

Having numerically characterized the performance of the various non-preconditioned PIMD integrators for the case of the harmonic oscillator external potential in Fig. 2.1, we now turn our attention to anharmonic external potentials. In this section, we consider both a weakly anharmonic (aHO) potential

$$V(q) = \Lambda\left(\frac{1}{2}q^2 + \frac{1}{10}q^3 + \frac{1}{100}q^4\right) \tag{2.46}$$

and the more strongly anharmonic quartic potential

$$V(q) = \frac{1}{4}q^4 \, . \tag{2.47}$$

All calculations are performed using $\hbar = 1$, $m = 1$ and $\beta = 1$. Assuming the system to be at room temperature (300 K), the thermal timescale corresponds to $\beta\hbar \approx 25.5$ fs and $\Lambda = m\omega^2$, where $\omega = 3315$ cm^{-1} for $\Lambda = 256$. The trajectories are performed with the centroid mode uncoupled from the thermostat (i.e., in the manner of T-RPMD); for the remaining $n-1$ internal modes, simulations performed with the OBABO and BAOAB schemes use the standard [18, 20] damping schedule of $\mathbf{\Gamma} = \mathbf{\Omega}$, and simulations performed using the Cayley modification (i.e., BCOCB, OBCBO, OMCMO, and OmCmO) use friction $\gamma_j = \min\{\omega_{j,n}, 0.9\gamma_j^{\max}(\Lambda), 0.9\gamma_j^{\max}(0)\}$ for the jth mode, where $\gamma_j^{\max}(\Lambda)$ is the friction that saturates the inequality in Eq. (2.21); for the quartic potential, we set $\Lambda = 1$ in this calculation of $\gamma_j^{\max}$.

Figures 2.2a and 2.2b presents kinetic energy expectation values for the aHO potential corresponding to 3315 cm^{-1} at room temperature. For the primitive kinetic energy expectation value, the results obtained using the various integration schemes with time-steps of both 0.5 fs (panel a) and 1.0 fs (panel b) are consistent with the observations for the harmonic potential in Fig. 2.1; specifically, the integrators without dimensionality freedom (OBABO, BAOAB, and OBCBO) fail to converge with increasing bead number, while the mollified integrators (OMCMO and OmCmO) smoothly converge with increasing bead number and the partially mollified scheme (OmCmO) is consistently more accurate than OBCBO and OMCMO. However, it is also clear that BCOCB exhibits the best accuracy with increasing bead number, converging to the exact result without perceivable time-step error.

Figures 2.2c and 2.2d present the corresponding results for the virial kinetic energy

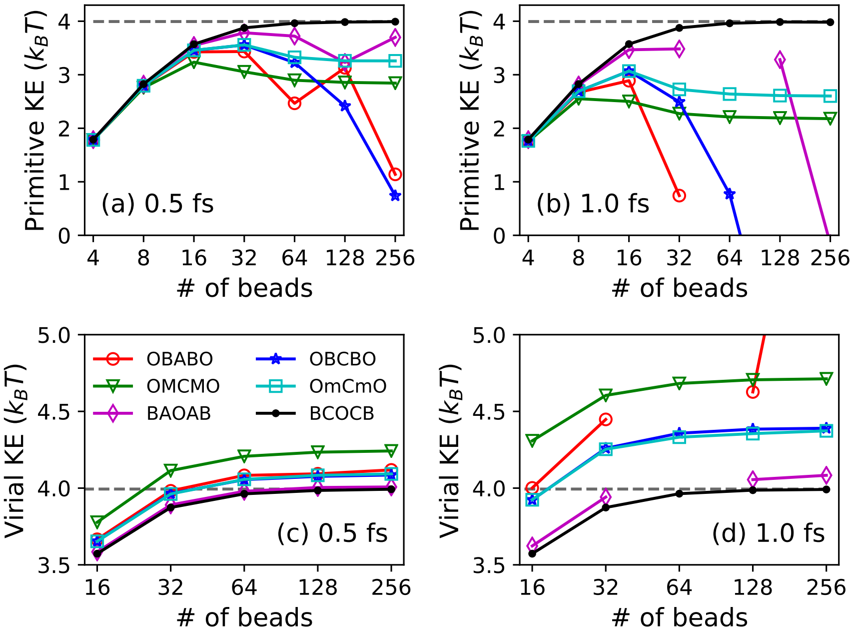

Figure 2.2. Primitive and virial kinetic energy expectation values as a function of bead number for the weakly anharmonic potential corresponding to 3315 cm^{-1} at room temperature, with results obtained using a time-step of 0.5 fs (panels a and c) and 1.0 fs (panels b and d). The standard error of all visible data points in each plot is smaller than the symbol size. The exact kinetic energy is indicated with a dashed line.

expectation value,

$$\langle \mathrm{KE}_n^{\mathrm{vir}} \rangle = \frac{1}{2\beta} - \frac{1}{2} \left\langle (\boldsymbol{q} - \bar{q}\mathbf{1}_n)^{\mathrm{T}} \boldsymbol{F}(\boldsymbol{q}) \right\rangle \tag{2.48}$$

where $\bar{q}$ is the centroid (bead-averaged) position and $\mathbf{1}_n$ is the n-dimensional vector with all entries equal to 1. Whereas the virial kinetic energy for all of the strongly stable integration schemes is well behaved, the OBABO and BAOAB schemes perform erratically at large time-steps due to their provable non-ergodicities [26]. Appealingly, the BCOCB scheme is consistently the most accurate for the virial kinetic energy expectation value, as it was for the primitive kinetic energy expectation value.

Figure 2.3a-d shows the results of the various numerical integration schemes for the primitive and virial kinetic energy expectation values, as a function of the MD time-step using 64 ring-polymer beads. Results are shown for both the aHO and the strongly anharmonic quartic oscillator. In all cases, the BCOCB scheme is consistently the most accurate across this array of model systems.

Finally, Fig. 2.3e illustrates the use of the BCOCB integrator for the calculation of real-time quantum dynamics within the T-RPMD approximation, where it replaces the often-employed OBABO integration scheme. Using 64 beads, the T-RPMD results are plotted for a range of integration time-steps. Strikingly, over the entire range of considered time-steps, BCOCB introduces negligible error in the calculated position time autocorrelation function; it is confirmed that these results are graphically indistinguishable from those obtained using the OBABO integrator in the small-time-step limit.

2.7 Results for liquid water

In model systems, Sec. 2.6 demonstrated the strong performance of the BCOCB integrator at evaluating both PIMD statistics and real-time dynamics through the T-RPMD model. In the current section, we test the accuracy and stability of the various un-mollified integration schemes (i.e., OBABO, OBCBO, BAOAB, and BCOCB) in liquid water, a high-dimensional and relatively complex system. Specifically, we consider a periodic 32-molecule water box at a temperature of 298 K and a density of 0.998 g/cm^3, as described by the q-TIP4P/F force field [41].

In Fig. 2.4, we compare the accuracy achieved by the different integrators for the average kinetic energy per hydrogen atom as a function of the number of ring-polymer beads. As in previous sections, we consider both the primitive (Eq. (2.29)) and virial (Eq. (2.48)) estimators for the kinetic energy. For each choice of integrator, time-step and bead number, the primitive and virial estimators for the kinetic energy of a single hydrogen atom were averaged over a 1-nanosecond trajectory integrated in the manner of T-RPMD (i.e., with the centroid mode uncoupled from the thermostat); for the remaining $n-1$ internal modes, simulations performed with the OBABO and BAOAB schemes use the standard [18, 20] damping schedule of $\boldsymbol{\Gamma} = \boldsymbol{\Omega}$, and simulations performed using the Cayley modification use friction $\gamma_j = \min\{\omega_j, 0.9\gamma_j^{\max}(0)\}$, where $\gamma_j^{\max}(\Lambda/m)$ saturates the inequality in Eq. (2.21) for the given values of j and Λ/m at the given time-step. Multi-nanosecond staging PIMD [8, 11] simulations at the small time-step of 0.1 fs were performed to obtain a bead-converged reference value for the H-atom kinetic energy, plotted as a dashed line in Figs. 2.4 and 2.5.

The primitive kinetic energy expectation values in Figs. 2.4a and 2.4b show similar trends to those seen in Figs. 2.1 and 2.2 for the harmonic and weakly anharmonic

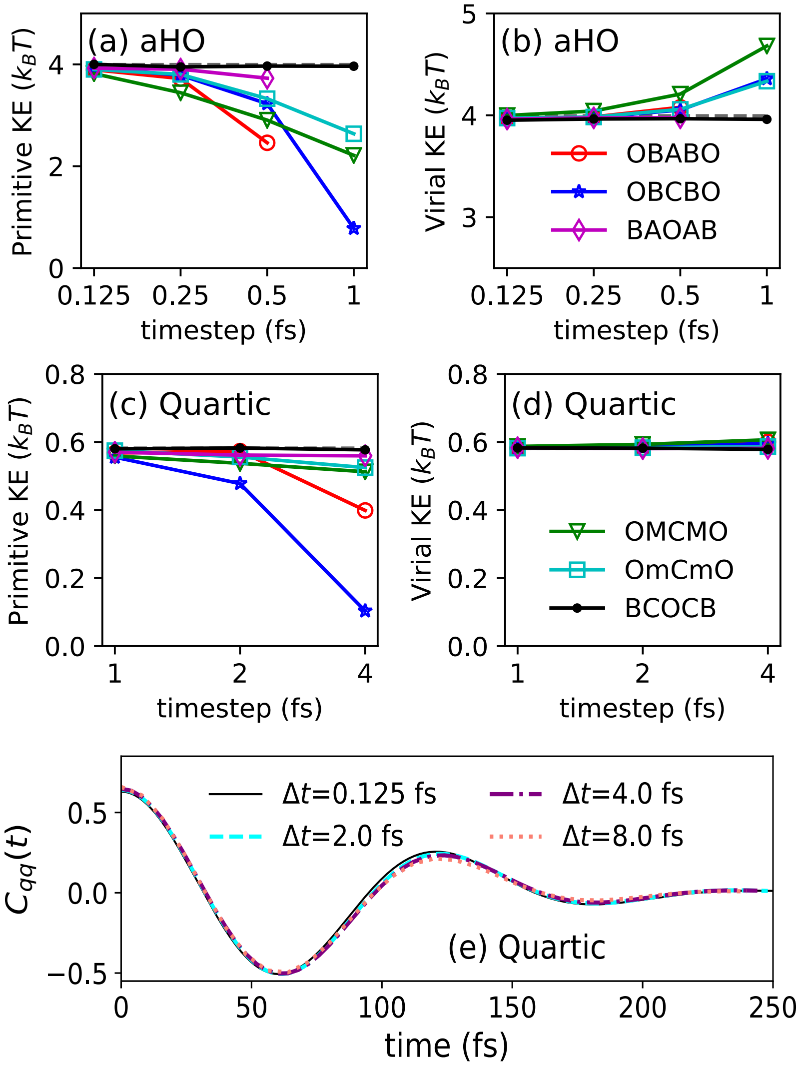

Figure 2.3. Primitive and virial kinetic energy expectation values as a function of the time-step for the weakly anharmonic potential corresponding to 3315 cm^{-1} at room temperature (panels a and b), and for the quartic potential (panels c and d). The exact kinetic energy is indicated with a dashed line. The standard error of all visible data points in each plot is smaller than the symbol size. Also, the position autocorrelation function (panel e) for the quartic oscillator at room temperature computed using T-RPMD with the BCOCB integrator. Results are obtained using 64 ring-polymer beads using time-steps of $\Delta t = 0.125, 2, 4,$ and 8 fs.

oscillators. For a 0.5-fs time-step (Fig. 2.4a), at which all integrators exhibit strong stability for ring polymers with up to 64 beads at the system temperature [26], the OBABO, BAOAB, and OBCBO primitive kinetic energy estimates diverge from the converged result as the number of beads increases, in agreement with the proven result that the error in the ring-polymer configurational distribution generated with these schemes grows unboundedly with increasing bead number. At the larger, 0.8-fs time-step, (Fig. 2.4b), OBABO and BAOAB formally lose strong stability and their respective primitive kinetic energy estimates dramatically diverge for bead numbers greater than 32; the strongly stable OBCBO scheme also yields a divergent result for the same reason as in Fig. 2.4a. As seen on the HO and aHO model systems, the primitive kinetic energy expectation value from the BCOCB integrator monotonically converges to the reference value with increasing bead number, avoiding any perceptible time-step error.

Figures 2.4c and 2.4d show the corresponding virial kinetic energy expectation values. For the smaller time-step of 0.5 fs, which is a common choice for path-integral simulations of water, all of the integrators perform similarly. However, upon increasing the time-step to 0.8 fs, significant differences in the performance of the integrators emerges, with only BCOCB avoiding perceptible time-step error.

To further compare the accuracy and stability of the OBABO, BAOAB, OBCBO, and BCOCB integrators, Fig. 2.5 considers the average kinetic energy per hydrogen atom obtained using 64 beads over a wide range of time-steps. These results show that BCOCB remains remarkably accurate for time-steps as large as 1.4 fs for liquid water, which corresponds to the limit of stability for Verlet integration of the centroid mode. In comparison, OBCBO diverges monotonically as the time-step increases, reaching unphysical values for the primitive expectation value and yielding sizable error (20%) for the virial expectation value. The erratic performance of both OBABO and BAOAB is due to the emergence of numerical resonance instabilities at time-steps greater than 0.6 fs at 64 beads; indeed, the largest safe time-step at which OBABO and BAOAB remain strongly stable at this bead number, $\Delta t_\star \approx 0.63$ fs, precedes the range of time-steps in Fig. 2.5 where these two integrators vary erratically.

Extending beyond statistics, we now consider quantum dynamical properties of liquid water within the T-RPMD approximation. Given the outstanding performance of the BCOCB scheme for the calculated statistical properties in Figs. 2.4 and 2.5, we present results that focus on this scheme in comparison to the most widely used

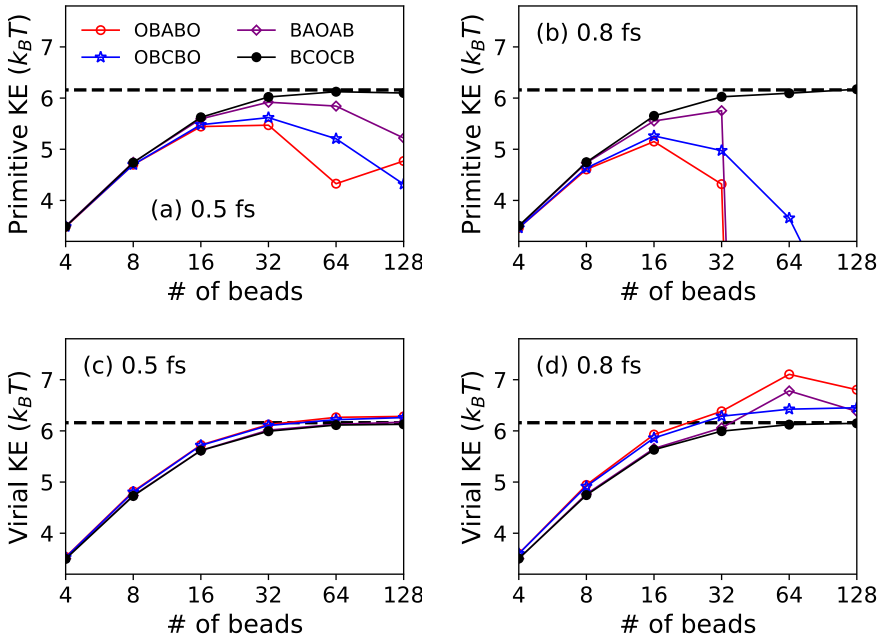

Figure 2.4. Primitive and virial kinetic energy expectation values as a function of the bead number for the hydrogen atom in liquid water at 298 K and 0.998 g/cm^3 at time-step $\Delta t = 0.5$ fs (panels a and c) and $\Delta t = 0.8$ fs (panels b and d). The reference kinetic energy, converged along a staging PIMD simulation at time-step $\Delta t = 0.1$ fs and 256 beads, is indicated with a dashed line. The standard error of all visible data points in each plot is smaller than the symbol size.

OBABO scheme. In particular, we consider the liquid water infrared absorption spectrum [42], which is proportional to $\omega^2 \tilde{I}(\omega)$ where the dipole spectrum $\tilde{I}(\omega) = \int_{\mathbb{R}} \mathrm{d}t\, \mathrm{e}^{-i\omega t} \tilde{C}_{\mu\cdot\mu}(t)$ is the Fourier transform of the Kubo-transformed dipole autocorrelation function $\tilde{C}_{\mu\cdot\mu}(t)$. The latter is approximated in the RPMD model by [16] $\tilde{C}_{\mu\cdot\mu}(t) = \frac{1}{N}\sum_{i=1}^{N} \langle \bar{\mu}_i(t) \cdot \bar{\mu}_i(0) \rangle$, where N is the number of molecules in the liquid, $\bar{\mu}_i(t)$ is the bead-averaged dipole moment of molecule i at time t and the angle brackets denote an average against the ring-polymer thermal distribution. To obtain the time-correlation functions and spectra shown in Fig. 2.6 for the OBABO and BCOCB integration schemes, 12-nanosecond T-RPMD trajectories were simulated for a ring-polymer with 64 beads and time-steps ranging from 0.2 to 1.4 fs, using the same friction schedule as described for Figs. 2.4 and 2.5.

Along each trajectory, the velocities of all degrees of freedom in the system were drawn anew from the Maxwell–Boltzmann distribution every 20 ps; the autocorrelation function was evaluated out to 2 ps by averaging over staggered windows of that

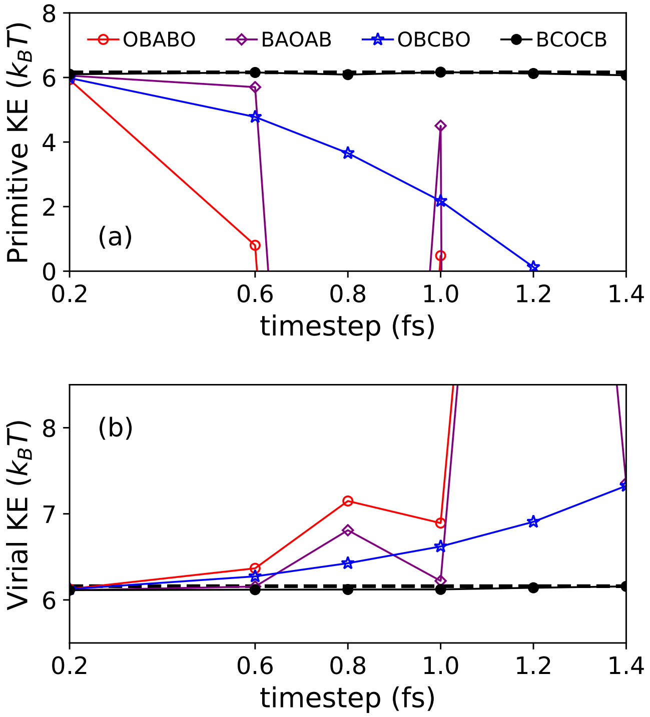

Figure 2.5. Primitive and virial kinetic energy expectation values as a function of the time-step for the hydrogen atom in liquid water at 298-K and 0.998 g/cm^3 as described by a 64-bead ring polymer. The reference kinetic energy, converged along a staging PIMD simulation at time-step $\Delta t = 0.1$ fs and 256 beads, is indicated with a dashed line. The standard error of all visible data points in each plot is smaller than the symbol size.

time-length within every 20-picosecond trajectory segment; and exponential-decay extrapolation was used to extend the autocorrelation function before evaluating its numerical Fourier transform to obtain the infrared absorption spectrum.

Figures 2.6a and 2.6b present the dipole autocorrelation functions obtained using the OBABO and BCOCB integrators with a range of time-steps. For the OBABO integrator, the calculated correlation function is qualitatively incorrect for time-steps as large as 0.8 fs. For the BCOCB integrator, the resulting correlations functions are far more robust with respect to time-step. Although modest differences are seen in the exponential tail of the correlation function, the dynamics on vibrational timescales (see inset) is largely unchanged as the time-step is varied from 0.2 fs to 1.4 fs. Figure 2.6c further emphasizes this point by showing the absorption spectrum that is obtained from the BCOCB time-correlation functions with the various time-steps. To minimize bias, we avoided any smoothing of the spectra shown in panel c. It is clearly seen that the librational and bending features (below 2500 cm^{-1}) are graphically indistinguishable over the entire range of considered time-steps. To clarify the comparison for the stretching region above 3000 cm^{-1}, we smooth the raw spectra in that region by convolution against a Gaussian kernel with a width of 150 cm^{-1} (see inset). Again, the robustness of the simulated spectrum over this span of time-steps is excellent, with the only significant effect due to finite-time-step error being a slight blue-shifting of the OH stretching frequency for the results using a 1.4-fs time-step, which is nearly three times larger than the typical value employed for the OBABO scheme for simulations with 64 beads. Taken together, these results indicate that the BCOCB integrator provides an excellent description of both PIMD statistics and T-RPMD dynamics in realistic molecular systems, substantially improving the accuracy and stability of previously employed numerical integrators.

2.8 Summary

In previous work [26], some of us showed that several widely used schemes for the non-preconditioned equations of motion of PIMD, including the widely used OBABO scheme, lack strong stability due to the use of exact free ring-polymer time evolution in the "A" sub-step, and we proved that this lack of strong stability gives rise to a lack of ergodicity in the thermostatted trajectories. We further showed that ergodicity can be restored by simply replacing the "A" sub-step with the Cayley transform.

In the current chapter, we show that a completely distinct, yet equally important, pathology exists in the "B" sub-step of previously developed non-preconditioned PIMD integrators, due to the outsized effect of the external potential on the dynamics of the high-frequency ring-polymer modes. Specifically, we show that previous integrators (including OBABO, BAOAB, and OBCBO) yield a numerical stationary distribution for which the overlap with the exact stationary distribution vanishes in the infinite-bead limit. We then show that this pathology is completely avoided in the BCOCB scheme, and we further show that the pathology can be eliminated in the OBCBO scheme by suitably mollifying the "B" sub-step, yielding the dimension-free non-preconditioned PIMD integrators, namely BCOCB, OMCMO, and OmCmO. Implementation of the dimension-free integration schemes involves no significant additional computational cost, no additional parameters, and no increase in algorithmic complexity in comparison to either OBABO or BAOAB. Furthermore, since the integrators considered here are all non-preconditioned, they can immedi-

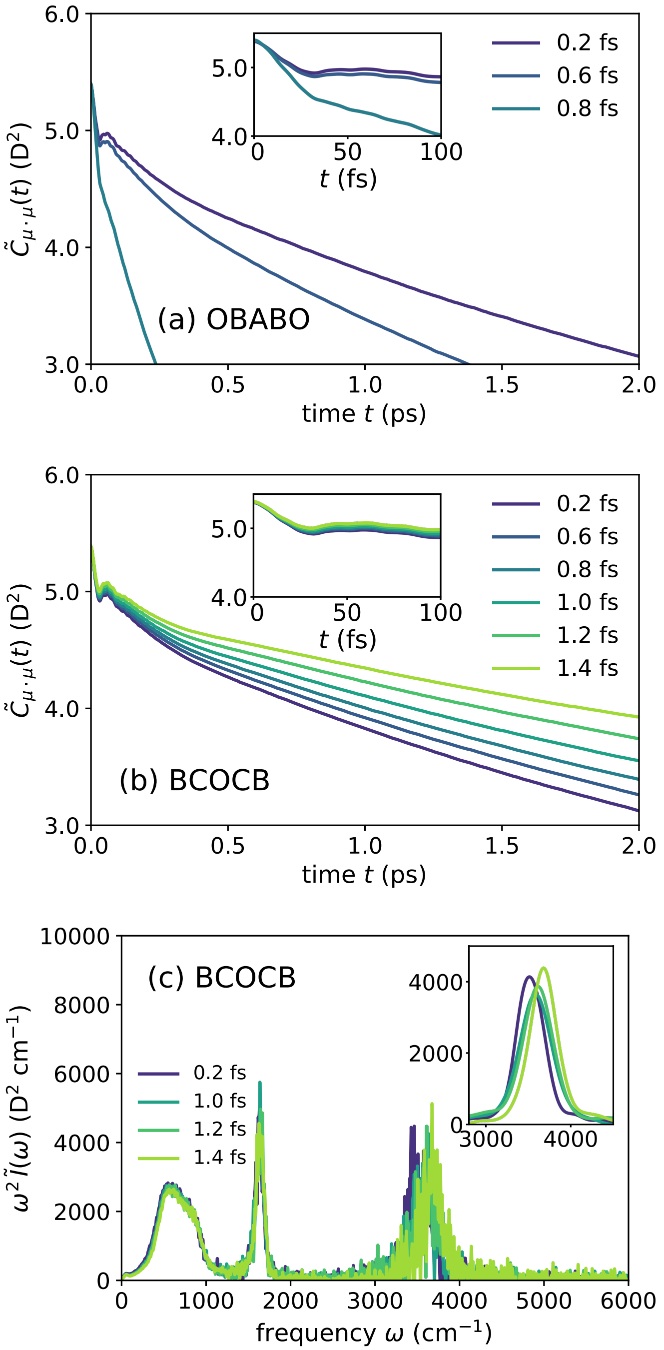

Figure 2.6. Dynamical properties of liquid water computed using T-RPMD with the (panel a) OBABO and (panels b and c) BCOCB integration schemes. Panels a and b present the Kubo-transformed dipole autocorrelation function computed with various time-steps, and panel c presents the absorption spectrum from the BCOCB correlation function at each time-step. The inset to panel c presents the OH stretching region with smoothing.

ately be used for computing the equilibrium statistical properties as well as dynamical properties within the T-RPMD approximation. The numerical performance of the BCOCB scheme is particularly striking, yielding results that are markedly better in terms of accuracy and time-step stability than any of the other considered integrators. For liquid water, it is shown that BCOCB allows for time-steps as large as 1.4 fs while exhibiting minimal time-step error in the calculation of both equilibrium expectation values and the dipole absorption spectrum.

Acknowledgments. We thank Andreas Eberle and Ondrej Marsalek for helpful discussions. N. B.-R. acknowledges support from the National Science Foundation under Award No. DMS-1816378 and the Alexander von Humboldt foundation. R. K., J. L. R.-R., and T. F. M. acknowledge support from the Department of Energy under Award No. DE-FOA-0001912 and the Office of Naval Research under Award No. N00014-10-1-0884.

2.9 Appendix A: Other splittings

There are exactly four locally third-order accurate symmetric splitting schemes that involve one new force evaluation per integration step and that involve splitting the T-RPMD dynamics into Hamiltonian and thermostat sub-steps: OBABO, BAOAB, OABAO, and ABOBA. In Sec. 2.3, we quantified the properties of OBABO, BAOAB, and their Cayley modifications in the case of a harmonic external potential. The corresponding properties of the Cayley modifications of OABAO and ABOBA are given below.

- OCBCO is exact in the velocity marginal, but the variance in the position marginal is $(\beta m_n)^{-1} s_{j,n,\Delta t}^2$ where $s_{j,n,\Delta t}^2 = (4m - \Delta t^2 \Lambda)/(4\Lambda + 4m\omega_{j,n}^2)$;
- CBOBC is exact in the position marginal, but the variance in the velocity marginal is $(\beta m_n)^{-1} r_{j,n,\Delta t}^2$ where $r_{j,n,\Delta t}^2 = 4m/(4m - \Delta t^2 \Lambda)$.

Numerical experiments confirmed these properties but did not show significant improvement in accuracy compared with BCOCB. Therefore, we did not include numerical results for these schemes.

Chapter 3

Generalization and optimization of dimension-free integrators for thermostatted ring-polymer molecular dynamics

Adapted from

J. L. Rosa-Raíces, J. Sun, N. Bou-Rabee, and T. F. Miller III, "A generalized class of strongly stable and dimension-free T-RPMD integrators", The Journal of Chemical Physics **154**, 024106 (2021) DOI: `10.1063/5.0036954`.

Abstract

Recent work, some presented in Chapter 2, showed that strong stability and dimensionality freedom are essential for robust numerical integration of thermostatted ring-polymer molecular dynamics (T-RPMD) and path-integral molecular dynamics (PIMD), without which standard integrators exhibit non-ergodicity and other pathologies [*J. Chem. Phys.* **151**, 124103 (2019); *J. Chem. Phys.* **152**, 104102 (2020)]. In particular, the BCOCB scheme, obtained via Cayley modification of the standard BAOAB scheme, features a simple reparametrization of the free ring-polymer sub-step that confers strong stability and dimensionality freedom and has been shown to yield excellent numerical accuracy in condensed-phase systems with large time-steps. Here, we build upon these findings by benchmarking the performance of BCOCB against a broad class of numerical integrators that also exhibit strong stability and dimensionality freedom for commonly used T-RPMD friction schedules. In addition to considering equilibrium accuracy and time-step stability as in Chapter 2, we compare the integrators on the basis of their rates of convergence to equilibrium and their efficiency at evaluating equilibrium expectation values. We find BCOCB to be optimal within this generalized class with respect to asymptotic statistical accuracy and efficiency for various configuration-dependent observables, although other integrators within the generalized class perform better for the evaluation of velocity-dependent quantities and exhibit faster convergence to equilibrium

at large bead numbers. Numerical evidence indicates that this result, here proven analytically for the harmonic oscillator, also holds for a strongly anharmonic model of liquid water. Both analytical and numerical results indicate that BCOCB outperforms other strongly stable and dimension-free T-RPMD integrators in terms of configurational accuracy, asymptotic sampling efficiency, and time-step stability across both toy-model and real-world applications.

3.1 Introduction

Path-integral molecular dynamics (PIMD) provides a practical and popular tool to simulate condensed-phase systems subject to strong nuclear quantum effects [1, 13, 43]. Based on the ring-polymer correspondence between quantum and classical Boltzmann–Gibbs statistics [14, 15], PIMD exploits the computational methods of molecular dynamics [44–47] to approximate quantum thermodynamics and kinetics through various classical models [12, 48–52]. Applications of PIMD include calculations of chemical reaction rates [53, 54], diffusion coefficients [16, 17], absorption spectra [42, 55], solid and liquid structure [56, 57], and equilibrium isotope effects [58, 59].

Many numerical integration schemes for PIMD are based on a symmetric Trotter splitting [27, 60] of the exact time-evolution operator, and feature a sub-step for free ring-polymer propagation [8, 11, 18]. Due to fast harmonic motions present in the free ring polymer, a *strongly stable* implementation of this sub-step is essential [61, 62]. Strong stability can be achieved by one of two approaches. The first approach introduces a preconditioned form of the equations of motion by modifying the ring-polymer mass matrix. Preconditioning improves the stability of the exact free ring-polymer update at the expense of consistent dynamics [3, 6, 8–11]. The second approach does not modify the ring-polymer mass matrix, leaving the dynamics non-preconditioned [7, 18–21], and instead replaces the exact free ring-polymer update with a strongly stable approximation [26]. We apply the latter approach in the current chapter to thermostatted ring-polymer molecular dynamics (T-RPMD) [20], a non-preconditioned variant of PIMD featuring an Ornstein–Uhlenbeck thermostat that approximately preserves the real-time dynamical accuracy of RPMD for quantum correlation functions of a wide range of observables [63].

In addition to strong stability of the free ring-polymer update, another basic requirement of a numerical integrator for T-RPMD is non-zero overlap between the numerically sampled and exact ring-polymer configurational distributions in the limit of an infinite number of ring-polymer beads. As argued in Chapter 2, standard integrators fail to satisfy this requirement at any finite integration time-step, which motivates the introduction of *dimension-free* T-RPMD schemes that allow for accurate configurational sampling with large time-stepping and arbitrarily many ring-polymer beads. We showed in Chapter 2 that standard integrators can be made dimension-free through the introduction of a suitable strongly stable ring-polymer update, and the current chapter investigates this finding in much greater generality.

To this end, we introduce a function θ that defines the free ring-polymer update and deduce how the choice of θ impacts the properties and performance of the corresponding T-RPMD integrator. The case $\theta(x) = x$, i.e., θ is the identity, corresponds to the exact free ring-polymer update. Therefore, to ensure second-order global accuracy, θ must approximate the identity near the origin, i.e., $\theta(0) = 0$, $\theta'(0) = 1$ and $\theta''(0) = 0$. Moreover, strong stability requires that the range of the function θ is within $(0, \pi)$ for $x > 0$, and ergodicity and dimensionality freedom of the corresponding T-RPMD integrator impose additional requirements on θ. There are many choices of θ that fulfill the identified conditions, including $\theta(x) = 2\arctan(x/2)$ which leads to the BCOCB scheme introduced in Chapter 2. In fact, we find that this choice of θ may be optimal for the estimation of configurational averages via T-RPMD from the perspectives of accuracy and efficiency in general applications.

This chapter is organized as follows. In Sec. 3.2 we recall T-RPMD and its time discretization, present the new function θ that determines the free ring-polymer update, and obtain sufficient conditions on θ to guarantee strong stability and dimensionality freedom of the corresponding T-RPMD integrator. In Sec. 3.3, we compare the performance of various θ in applications to the one-dimensional quantum harmonic oscillator and to a quantum-mechanical model of room-temperature liquid water. Section 3.4 summarizes the chapter, and subsequent appendices provide supporting mathematical proofs and list calculation details.

3.2 Theory

3.2.1 T-RPMD integrator desiderata

We review the setting addressed by T-RPMD integrators, recalling details from Chapter 2 as necessary. Consider a one-dimensional molecular system with potential energy function $V(q)$, mass m, and Hamiltonian $\hat{H}$ as defined in Eq. (2.2). Following the arguments and definitions in Sec. 2.2, the equations of motion for the corresponding n-bead ring polymer held at constant temperature T by an Ornstein–Uhlenbeck thermostat are

$$\begin{aligned} \dot{\boldsymbol{q}}(t) \;=\; \boldsymbol{v}(t) \ , \ \dot{\boldsymbol{v}}(t) \;=\; & -\boldsymbol{\Omega}^2\boldsymbol{q}(t) + m_n^{-1}\boldsymbol{F}(\boldsymbol{q}(t)) \\ & -\boldsymbol{\Gamma}\boldsymbol{v}(t) + \sqrt{2\beta^{-1}m_n^{-1}}\,\boldsymbol{\Gamma}^{1/2}\,\dot{\boldsymbol{W}}(t) \ , \end{aligned} \tag{3.1}$$

which reproduces Eq. (2.8). Here, $\boldsymbol{W}$ is an n-dimensional standard Brownian motion, $\boldsymbol{q}(t) = (q_0(t), \ldots, q_{n-1}(t))^{\mathrm{T}}$ is the vector of positions for the n ring-polymer beads at time $t \geq 0$, $\boldsymbol{v}(t)$ are the corresponding velocities, $m_n = m/n$, $\beta = (k_B T)^{-1}$, and $\boldsymbol{F}(\boldsymbol{q}) = -\nabla V_n^{\mathrm{ext}}(\boldsymbol{q})$ where V_n^{ext} is the bead-averaged potential defined in (2.6). Moreover, $\boldsymbol{\Omega}^2$ is the $n \times n$ symmetric positive semi-definite matrix given in Eq. (2.10), diagonalizable by an $n \times n$ orthonormal real discrete Fourier transform matrix $\boldsymbol{U}$ as

$$\boldsymbol{\Omega} = \boldsymbol{U}\,\mathrm{diag}\,(0, \omega_{1,n}, \ldots, \omega_{n-1,n})\,\boldsymbol{U}^{\mathrm{T}} \ , \tag{3.2}$$

where $\omega_{j,n}$ is the jth Matsubara frequency defined in (2.11). Similarly, $\boldsymbol{\Gamma}$ in Eq. (2.8) is an $n \times n$ symmetric positive semi-definite friction matrix of the form

$$\boldsymbol{\Gamma} = \boldsymbol{U}\,\mathrm{diag}\,(0, \gamma_{1,n}, \ldots, \gamma_{n-1,n})\,\boldsymbol{U}^{\mathrm{T}} \ , \tag{3.3}$$

where $\gamma_{j,n}$ is the friction factor in the jth normal mode.

In RPMD and T-RPMD calculations, one is often interested in the dynamics of Eq. (3.1) with initial conditions drawn from the stationary distribution with non-normalized density $\exp(-\beta H_n(\boldsymbol{q}, \boldsymbol{v}))$, where $H_n(\boldsymbol{q}, \boldsymbol{v})$ is the ring-polymer Hamiltonian defined by Eqs. (2.5)–(2.7). We shall consider this case, as well as the case of initialization from non-equilibrium distributions, later in this chapter.

Numerical integrators for Eq. (3.1) typically employ symmetric propagator splittings of the form given in Eq. (2.1), where the operator $\mathcal{L}_n = \mathcal{A}_n + \mathcal{B}_n + \mathcal{O}_n$ includes contributions from the n-bead free ring-polymer motion ($\mathcal{A}_n$), the external potential ($\mathcal{B}_n$) and the thermostat ($\mathcal{O}_n$), and Δt is a sufficiently small time-step. Recall from Chapter 2 that the standard RPMD integrator is recovered in the limit of zero coupling to the thermostat [18], and that Eq. (2.1) yields the OBABO scheme of Bussi and Parrinello [22] if $a = 1$ and the BAOAB scheme of Leimkuhler [25] if $a = 0$. Standard implementations of the T-RPMD splittings in Eq. (2.1) use the exact free ring-polymer propagator $\exp((\Delta t/2)\mathcal{A}_n)$ to evolve the free ring polymer; however, previous work by some of us [26] showed that such implementations exhibit poor ergodicity if large numbers n of ring-polymer beads are employed in conjunction with large time-steps Δt, and suggested replacing the exact ring-polymer propagator with its Cayley approximation [64] for improved performance. Subsequent work, presented in Chapter 2, introduced a Cayley-modified BAOAB scheme, denoted BCOCB, and presented numerical evidence that cemented the scheme as an improvement over standard BAOAB due to its superior equilibrium accuracy and time-step stability.

Generalizing beyond the Cayley modification, the current chapter studies a family of modified BAOAB schemes that contains BCOCB and introduces others with similar theoretical guarantees. These BAOAB-like schemes are obtained by replacing the exact free ring-polymer update in Eq. (2.1) with approximations that endow the properties listed below.

- (**P1**) *Strong stability.* For a free ring polymer (i.e., for $V \equiv \text{const.}$), the integrator with $\gamma_{j,n} = 0$ is both strongly stable and second-order accurate in Δt.
- (**P2**) *Free ring-polymer ergodicity.* For a free ring polymer, the integrator with $\gamma_{j,n} > 0$ is ergodic with respect to the distribution with density proportional to $\exp(-\beta H_n^0(\boldsymbol{q}, \boldsymbol{v}))$.
- (**P3**) *Dimension-free stability.* For a harmonically confined ring polymer (i.e., for $V(q) = (\Lambda/2)\, q^2$), the integrator with $\gamma_{j,n} = 0$ is stable for any n if Δt leads to stable integration for $n = 1$.
- (**P4**) *Dimension-free ergodicity.* For a harmonically confined ring polymer, the integrator with $\gamma_{j,n} > 0$ and stable Δt is ergodic with respect to its stationary distribution for any n.
- (**P5**) *Dimension-free equilibrium accuracy.* For a harmonically confined ring polymer, the integrator leaves invariant an accurate approximation of the distribution with density proportional to $\exp\big(-(\beta m_n/2)\, \boldsymbol{q}^{\mathrm{T}}(\Lambda/m + \boldsymbol{\Omega}^2)\boldsymbol{q}\big)$, with bounded error for any n.

To obtain integrators satisfying properties (**P1**)–(**P5**), we introduce a function θ that defines the free ring-polymer update and then place conditions on θ that lead to the desired properties. To this end, let

$$\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2} = \boldsymbol{\mathcal{Q}}_{j,n} \begin{bmatrix} \mathrm{e}^{i\theta(\omega_{j,n}\Delta t)/2} & 0 \\ 0 & \mathrm{e}^{-i\theta(\omega_{j,n}\Delta t)/2} \end{bmatrix} \boldsymbol{\mathcal{Q}}_{j,n}^{-1} , \tag{3.4}$$

where $\boldsymbol{\mathcal{Q}}_{j,n} = \left[\begin{smallmatrix} 1 & 1 \\ i\omega_{j,n} & -i\omega_{j,n} \end{smallmatrix}\right]$ and essential properties of θ are determined in the sequel. We focus on T-RPMD schemes derived from the BAOAB splitting (i.e., $a = 0$ in Eq. (2.1)) with the exact free ring-polymer update replaced by $\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2}$. For such schemes, an integration time-step is comprised by the following sequence of sub-steps:

B: Update velocities for half a step: $\boldsymbol{v} \leftarrow \boldsymbol{v} + \frac{\Delta t}{2} m_n^{-1} \boldsymbol{F}(\boldsymbol{q})$.

Convert bead Cartesian coordinates to normal modes using

$$\boldsymbol{\varrho} = \boldsymbol{U}^{\mathrm{T}} \boldsymbol{q} \quad \text{and} \quad \boldsymbol{\varphi} = \boldsymbol{U}^{\mathrm{T}} \boldsymbol{v} . \tag{3.5}$$

A: Evolve the free ring polymer in normal-mode coordinates for half a step:

$$\begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \leftarrow \boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2} \begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \quad \text{for} \quad 0 \le j \le n-1 .$$

O: Perform an Ornstein–Uhlenbeck velocity update for a full time-step:

$$\varphi_j \leftarrow \mathrm{e}^{-\gamma_{j,n}\Delta t} \varphi_j + \sqrt{\frac{1 - \mathrm{e}^{-2\gamma_{j,n}\Delta t}}{\beta m_n}} \xi_j ,$$

where ξ_j are independent standard normal random variables and $0 \le j \le n-1$.

A: Evolve the free ring polymer in normal-mode coordinates for half a step:

$$\begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \leftarrow \boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2} \begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \quad \text{for} \quad 0 \le j \le n-1 .$$

Convert back to bead Cartesian coordinates using the inverse of $\boldsymbol{U}$, which is just its transpose since $\boldsymbol{U}$ is orthogonal.

B: Update velocities for half a step: $\boldsymbol{v} \leftarrow \boldsymbol{v} + \frac{\Delta t}{2} m_n^{-1} \boldsymbol{F}(\boldsymbol{q})$.

In the remainder of this section, we identify conditions on the choice of θ that imply properties (**P1**)–(**P5**) for the corresponding T-RPMD integrator. Despite our focus on BAOAB-like splittings, we describe how the conditions on θ can be adjusted to construct integrators derived from the OBABO splitting (i.e., $a = 1$ in Eq. (2.1)) that satisfy properties (**P1**)–(**P5**).

3.2.2 Strong stability of RPMD with a constant external potential

In this section, sufficient conditions on θ are identified to satisfy property (**P1**) in Sec. 3.2.1. Let $V \equiv$ const. and $\gamma_{j,n} = 0$ for $1 \le j \le n-1$, corresponding to the free ring polymer. The jth normal mode $[\varrho_j, \varphi_j]^{\mathrm{T}}$ satisfies

$$\begin{bmatrix} \dot{\varrho}_j \\ \dot{\varphi}_j \end{bmatrix} = \boldsymbol{A}_{j,n} \begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \quad \text{where} \quad \boldsymbol{A}_{j,n} = \begin{bmatrix} 0 & 1 \\ -\omega_{j,n}^2 & 0 \end{bmatrix} . \tag{3.6}$$

In this case, the algorithm from Sec. 3.2.1 reduces to a step of $\boldsymbol{\mathcal{S}}_{j,n,\Delta t} \approx \exp(\Delta t \boldsymbol{A}_{j,n})$, i.e.,

$$\begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \leftarrow \boldsymbol{\mathcal{S}}_{j,n,\Delta t} \begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \quad \text{for} \ \ 0 \le j \le n-1 \, , \tag{3.7}$$

where $\boldsymbol{\mathcal{S}}_{j,n,\Delta t} = \boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t} \boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t}$ follows from Eq. (3.4) and the function θ is such that property (**P1**) holds.

We proceed to identify sufficient conditions on θ such that the corresponding free ring-polymer update satisfies property (**P1**). First note that for any function θ such that $\theta(-x) = -\theta(x)$ for $x > 0$, the structure of $\boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t}$ guarantees that the corresponding free ring-polymer update is reversible, symplectic, and preserves the free ring-polymer Hamiltonian $H_n^0(\boldsymbol{q}, \boldsymbol{v})$. Now, observe that $\boldsymbol{\mathcal{S}}_{j,n,\Delta t}$ is exact if $\theta(x) = x$; therefore, second-order accuracy requires that θ approximates the identity near the origin, i.e.,

$$\theta(0) = 0, \ \theta'(0) = 1, \ \text{and} \ \ \theta''(0) = 0. \tag{\textbf{C1}}$$

Moreover, strong stability follows if the eigenvalues $\exp(\pm i \theta(\omega_{j,n} \Delta t))$ of $\boldsymbol{\mathcal{S}}_{j,n,\Delta t}$ are distinct [26]; to this end we require that

$$0 < \theta(x) < \pi \ \ \text{for} \ \ x > 0 \, . \tag{\textbf{C2}}$$

Jointly, conditions (**C1**) and (**C2**) guarantee that the update in Eq. (3.7) satisfies property (**P1**). There are many different choices of θ that obey these conditions, e.g., $\theta(x) = \arctan(x)$, $\theta(x) = \arccos(\mathrm{sech}(x))$[1] and $\theta(x) = 2\arctan(x/2)$. The latter choice leads to the Cayley approximation of the free ring-polymer update, as can be verified by substitution in Eq. (3.4) and comparison of the resulting $\boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t}$ with Eq. (17) in Chapter 2. Figure 3.1 compares the eigenvalues of $\boldsymbol{\mathcal{S}}_{j,n,\Delta t}$ with $\theta(x) = x$ and several choices of θ that meet conditions (**C1**) and (**C2**).

---

[1] The function $\theta(x) = \arccos(\mathrm{sech}(x))$ is not differentiable at the origin and hence, strictly speaking, does not satisfy condition (**C1**). Moreover, the function has even symmetry and hence fails to yield a reversible free ring-polymer update. These formal shortcomings can be fixed by multiplying the function by $\mathrm{sign}(x)$, which we implicitly do for this and other functions θ with similar features.

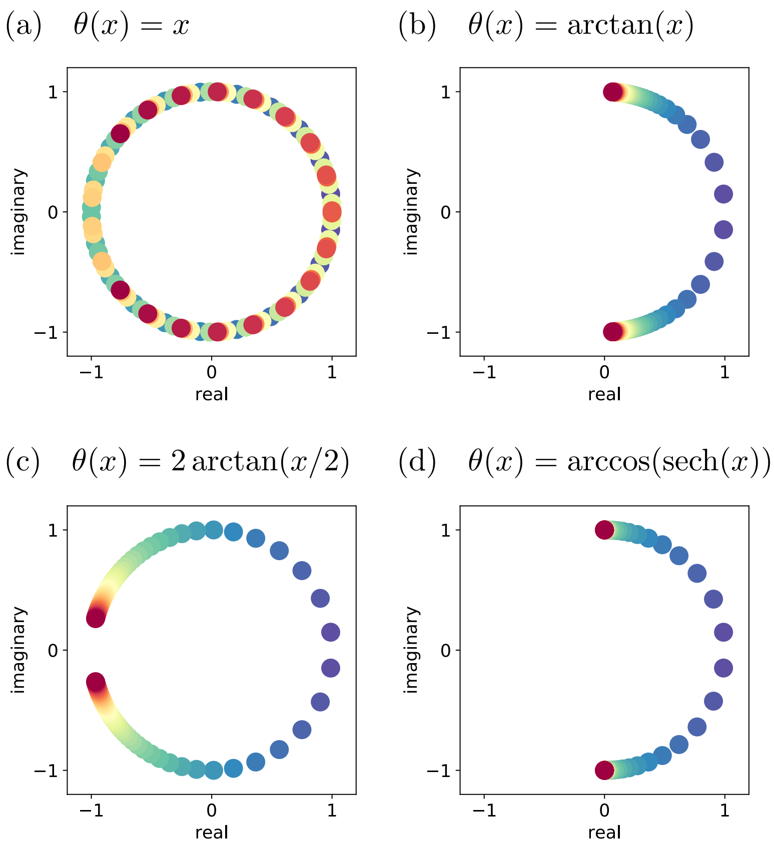

(a) $\theta(x) = x$

(b) $\theta(x) = \arctan(x)$

(c) $\theta(x) = 2\arctan(x/2)$

(d) $\theta(x) = \arccos(\operatorname{sech}(x))$

Figure 3.1. Eigenvalues of $\boldsymbol{\mathcal{S}}_{j,n,\Delta t}$ for 50 different time-steps between 0.05 and 5.0 (evenly spaced) and fixed Matsubara frequency $\omega = 3$. The colors go from blue (smallest time-step) through green and yellow to red (largest time-step). In panel a, the eigenvalues rotate around the unit circle several times, which indicates that the corresponding $\boldsymbol{\mathcal{S}}_{j,n,\Delta t}$ is not always strongly stable. In panels b, c, and d, the eigenvalues are distinct and on the unit circle; thus the corresponding $\boldsymbol{\mathcal{S}}_{j,n,\Delta t}$ is strongly stable.

3.2.3 Ergodicity of T-RPMD with a constant external potential

In this section, it is shown that condition (**C2**) implies property (**P2**) in Sec. 3.2.1. Let $V \equiv$ const. and $\gamma_{j,n} > 0$ for $1 \leq j \leq n-1$, corresponding to the free ring polymer with a Ornstein–Uhlenbeck thermostat. In this case, the jth normal mode satisfies

$$\begin{bmatrix} \dot{\varrho}_j \\ \dot{\varphi}_j \end{bmatrix} = (\boldsymbol{A}_{j,n} + \boldsymbol{O}_{j,n}) \begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} + \begin{bmatrix} 0 \\ \sqrt{\frac{2\gamma_{j,n}}{\beta m_n}} \dot{W}_j \end{bmatrix} , \tag{3.8}$$

where $\boldsymbol{O}_{j,n} = \left[\begin{smallmatrix} 0 & 0 \\ 0 & -\gamma_{j,n} \end{smallmatrix}\right]$ and $\dot{W}_j$ is a scalar white-noise. The solution $[\varrho_j(t), \varphi_j(t)]^{\mathrm{T}}$ of Eq. (3.8) is an ergodic Markov process, and in the limit as $t \to \infty$, its distribution converges to the centered bivariate normal with covariance

$$\boldsymbol{\Sigma}_{j,n} = \frac{1}{\beta m_n}\begin{bmatrix} s_{j,n}^2 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{where} \quad s_{j,n}^2 = \frac{1}{\omega_{j,n}^2} \ .$$

This distribution corresponds to the jth marginal of the free ring-polymer equilibrium distribution with density proportional to $\exp(-\beta H_n^0(\boldsymbol{q},\boldsymbol{v}))$.

The choice of $\gamma_{j,n} > 0$ in Eq. (3.8) determines the rate at which the associated Markov process converges to its stationary distribution if initialized away from it. When $\gamma_{j,n} < 2\omega_{j,n}$, the process is dominated by the deterministic Hamiltonian dynamics and is characterized as *underdamped*; on the other hand, when $\gamma_{j,n} > 2\omega_{j,n}$, the process is *overdamped*; and at the critical value $\gamma_{j,n} = 2\omega_{j,n}$ the process is characterized as *critically damped* and converges to equilibrium fastest [65, 66]. This analytical result motivates the so-called PILE friction schedule [18, 20]. We specialize to this schedule in the remainder of this section and set $\gamma_{j,n} = 2\omega_{j,n}$ for $1 \le j \le n-1$.

The BAOAB-like update in Sec. 3.2.1, applied to Eq. (3.8) with PILE friction, can be written compactly as

$$\begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \leftarrow \boldsymbol{\mathcal{M}}_{j,n,\Delta t}\begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} + \boldsymbol{\mathcal{R}}_{j,n,\Delta t}^{1/2}\begin{bmatrix} \xi_j \\ \eta_j \end{bmatrix} \quad \text{for} \ \ 0 \le j \le n-1 \ , \tag{3.9}$$

where ξ_j and η_j are independent standard normal random variables and we have introduced the 2×2 matrices

$$\boldsymbol{\mathcal{M}}_{j,n,\Delta t} = \boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2}\boldsymbol{\mathcal{O}}_{j,n,\Delta t}\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2} \ , \quad \boldsymbol{\mathcal{O}}_{j,n,\Delta t} = \begin{bmatrix} 1 & 0 \\ 0 & \mathrm{e}^{-2\omega_{j,n}\Delta t} \end{bmatrix} \ \text{and}$$

$$\boldsymbol{\mathcal{R}}_{j,n,\Delta t} = \frac{1-\mathrm{e}^{-4\omega_{j,n}\Delta t}}{\beta m_n}\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2}\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}(\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2})^{\mathrm{T}} \ .$$

Since $\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2}$ and the Ornstein–Uhlenbeck update are individually preservative irrespective of the chosen θ, Eq. (3.9) exactly preserves the free ring-polymer equilibrium distribution for any choice of θ that satisfies (**C1**) and (**C2**).

The ergodicity of the integrator specified by Eq. (3.9) depends entirely on the asymptotic stability of $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$, that is, on whether $\|\boldsymbol{\mathcal{M}}_{j,n,\Delta t}^k\| \to 0$ as $k \to \infty$, where $\|\cdot\|$ is a matrix norm. The matrix $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$ is asymptotically stable if its *spectral radius* (i.e., the modulus of its largest eigenvalue) is smaller than unity [62], which depends on

$$\det(\boldsymbol{\mathcal{M}}_{j,n,\Delta t}) = \mathrm{e}^{-2\omega_{j,n}\Delta t} \ \text{ and } \ \mathrm{Tr}(\boldsymbol{\mathcal{M}}_{j,n,\Delta t}) = \cos(\theta(\omega_{j,n}\Delta t))(1+\mathrm{e}^{-2\omega_{j,n}\Delta t}) \ .$$

In particular, the eigenvalues of $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$ are both inside the unit circle if and only if

$$|\operatorname{Tr}(\boldsymbol{\mathcal{M}}_{j,n,\Delta t})| < 1 + \det(\boldsymbol{\mathcal{M}}_{j,n,\Delta t}) < 2\ ;$$

a proof of this claim is provided in Appendix 3.5. This inequality reveals that condition (**C2**) implies property (**P2**). Moreover, if $\operatorname{Tr}(\boldsymbol{\mathcal{M}}_{j,n,\Delta t})^2 - 4\det(\boldsymbol{\mathcal{M}}_{j,n,\Delta t}) \le 0$, then the spectral radius of $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$ is minimal and equal to $\sqrt{\det(\boldsymbol{\mathcal{M}}_{j,n,\Delta t})} = \mathrm{e}^{-\omega_{j,n}\Delta t}$; this occurs when $|\cos(\theta(\omega_{j,n}\Delta t))| \le \operatorname{sech}(\omega_{j,n}\Delta t)$ for all $\omega_{j,n}\Delta t$, which holds if the function θ satisfies

$$\arccos(\operatorname{sech}(x)) \le \theta(x) \le \pi - \arccos(\operatorname{sech}(x)) \quad \text{for} \quad x > 0\ . \tag{3.10}$$

Any choice of θ that does not satisfy Eq. (3.10) will be overdamped in some modes, in the sense that the corresponding $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$ will have a spectral radius strictly larger than $\mathrm{e}^{-\omega_{j,n}\Delta t}$.

The function $\theta(x) = \arccos(\operatorname{sech}(x))$ saturates the (left) inequality in Eq. (3.10) while satisfying conditions (**C1**) and (**C2**), and hence provides a strongly stable and critically damped integrator for the thermostatted free ring polymer. As illustration of this, Fig. 3.2a shows that $\theta(x) = \arctan(x)$ is overdamped for all modes whereas the Cayley angle $\theta(x) = 2\arctan(x/2)$ exhibits mixed damping. In contrast, the function $\theta(x) = \arccos(\operatorname{sech}(x))$ preserves the critically damped behavior of its continuous counterpart under the PILE friction schedule. Figure 3.2b confirms that the spectral radius of $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$ is minimal at $\theta(x) = \arccos(\operatorname{sech}(x))$ for $x > 0$; consequently, this choice of θ optimizes the convergence of the integrator to stationarity.

Conditions (**C1**) and (**C2**) also imply property (**P2**) for the OBABO-like update associated with a compliant choice of θ, because the matrices $\boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t}\boldsymbol{\mathcal{O}}_{j,n,\Delta t}\boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t}$ and $\boldsymbol{\mathcal{O}}^{1/2}_{j,n,\Delta t}\boldsymbol{\mathcal{S}}_{j,n,\Delta t}\boldsymbol{\mathcal{O}}^{1/2}_{j,n,\Delta t}$ have equal spectral radii.

3.2.4 Dimension-free stability of RPMD with a harmonic external potential

In this section, we identify a condition on θ that yields property (**P3**) in Sec. 3.2.1. Let $V(q) = (\Lambda/2)\,q^2$ and $\gamma_{j,n} = 0$ for $1 \le j \le n-1$, corresponding to the non-thermostatted ring polymer with a harmonic external potential. In this case, the jth normal mode satisfies

$$\begin{bmatrix} \dot{\varrho}_j \\ \dot{\varphi}_j \end{bmatrix} = (\boldsymbol{A}_{j,n} + \boldsymbol{B}) \begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \tag{3.11}$$

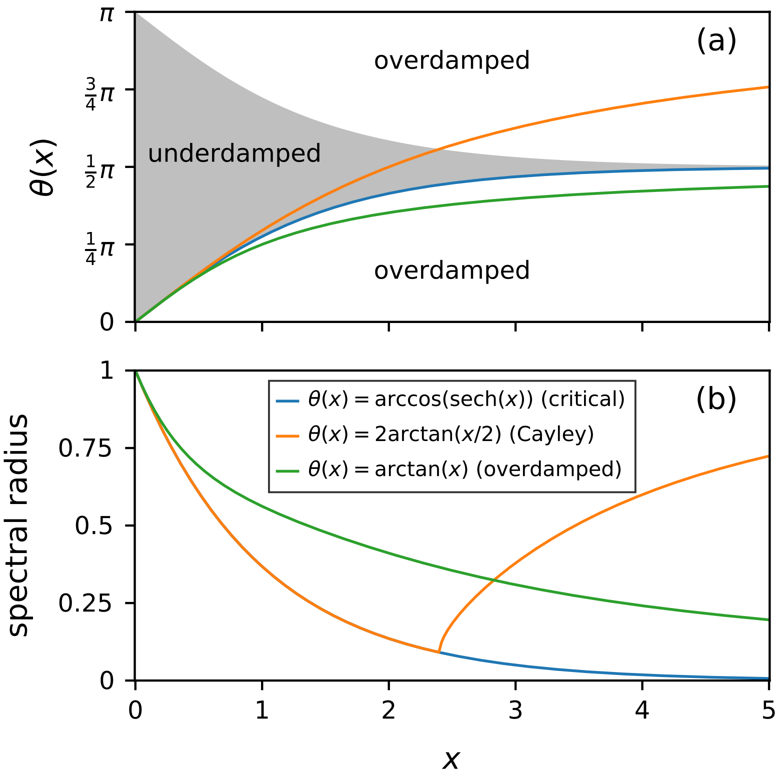

Figure 3.2. Spectral properties of the T-RPMD update for the free ring polymer for various choices of θ. Panel a plots the functions $\theta(x) = \arccos(\mathrm{sech}(x))$, $\arctan(x)$ and $\arctan(x/2)$, and regions of overdamping and underdamping with PILE friction, separated at the locus of points where $|\cos(\theta(x))\cosh(x)| = 1$. The dynamics is underdamped in the gray region ($|\cos(\theta(x))\cosh(x)| < 1$), while in the white region ($|\cos(\theta(x))\cosh(x)| > 1$) the dynamics is overdamped. The function $\theta(x) = \arctan(x)$ lies in the overdamped region for $x > 0$, whereas $\theta(x) = 2\arctan(x/2)$ is in the underdamped region for $x \lessapprox 2.4$ and in the overdamped region otherwise. The function $\theta(x) = \arccos(\mathrm{sech}(x))$, however, is critically damped for $x > 0$ and optimizes the convergence rate of the integrator. Panel b plots the spectral radius of $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$ corresponding to each choice of θ as a function of x.

where $\boldsymbol{B} = \left[\begin{smallmatrix} 0 & 0 \\ -\Lambda/m & 0 \end{smallmatrix}\right]$, and conserves the Hamiltonian

$$H_{j,n}(\rho_j, \varphi_j) = \frac{m_n}{2}\left(|\varphi_j|^2 + (\omega_{j,n}^2 + \Lambda/m)|\varrho_j|^2\right) .$$

For this system, the BAOAB-like update in Sec. 3.2.1 reduces to

$$\begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \leftarrow \boldsymbol{\mathcal{M}}_{j,n,\Delta t} \begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} \quad \text{for} \quad 0 \le j \le n-1 \, , \tag{3.12}$$

where we have introduced the 2×2 matrices

$$\boldsymbol{\mathcal{M}}_{j,n,\Delta t} = \boldsymbol{\mathcal{B}}^{1/2} \boldsymbol{\mathcal{S}}_{j,n,\Delta t} \boldsymbol{\mathcal{B}}^{1/2} \quad \text{and} \quad \boldsymbol{\mathcal{B}}^{1/2} = \begin{bmatrix} 1 & 0 \\ -\Delta t(\Lambda/m)/2 & 1 \end{bmatrix} \, .$$

This update may be interpreted as a symplectic perturbation of the free ring-polymer update in Eq. (3.7) due to the harmonic external potential [26], and conserves a modification of $H_{j,n}$ that depends on the choices of θ and Δt [67].

The update in Eq. (3.12) is stable if [9]

$$\max_{0\le j\le n-1} \frac{1}{2} |\operatorname{Tr}(\boldsymbol{\mathcal{M}}_{j,n,\Delta t})| = \max_{0\le j\le n-1} |\mathcal{A}_{j,n}| < 1 \; , \tag{3.13}$$

where

$$\mathcal{A}_{j,n} = \cos(\theta(\omega_{j,n}\Delta t)) - \frac{\Delta t^2(\Lambda/m)}{2} \frac{\sin(\theta(\omega_{j,n}\Delta t))}{\omega_{j,n}\Delta t} \; .$$

Moreover, the 0th (i.e., centroid) mode, like the single-bead ring polymer, evolves through the velocity Verlet algorithm, whose stability requires that $\Delta t^2\Lambda/m < 4$. Combining this requirement with condition (**C2**) yields a sufficient condition for Eq. (3.13) to hold at any n,

$$0 < \theta(x) \le 2\arctan(x/2) \quad \text{for} \quad x > 0 \; . \tag{\textbf{C3}}$$

A proof of this result is provided in Sec. 3.6. The functions $\theta(x) = 2\arctan(x/2)$, $\theta(x) = \arctan(x)$ and $\theta(x) = \arccos(\operatorname{sech}(x))$ all satisfy condition (**C3**), which ensures that the corresponding RPMD integrator meets property (**P3**).

3.2.5 Dimension-free ergodicity and equilibrium accuracy of T-RPMD with a harmonic external potential

In this section, it is shown that condition (**C3**) implies property (**P4**) in Sec. 3.2.1, and an additional condition is introduced to ensure that property (**P5**) holds. Let $V(q) = (\Lambda/2)\, q^2$ and $\gamma_{j,n} = 2\omega_{j,n}$ for $1 \le j \le n-1$. In this case, the jth normal mode satisfies

$$\begin{bmatrix} \dot{\varrho}_j \\ \dot{\varphi}_j \end{bmatrix} = (\boldsymbol{A}_{j,n} + \boldsymbol{B} + \boldsymbol{O}_{j,n}) \begin{bmatrix} \varrho_j \\ \varphi_j \end{bmatrix} + \begin{bmatrix} 0 \\ \sqrt{\frac{4\omega_{j,n}}{\beta m_n}}\dot{W}_j \end{bmatrix} \; . \tag{3.14}$$

The solution $[\varrho_j(t), \varphi_j(t)]^{\mathrm{T}}$ of Eq. (3.14) is an ergodic Markov process, and its distribution as $t \to \infty$ converges to the centered bivariate normal with covariance matrix

$$\boldsymbol{\Sigma}_{j,n} = \frac{1}{\beta m_n} \begin{bmatrix} s_{j,n}^2 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{where} \quad s_{j,n}^2 = \frac{1}{\Lambda/m + \omega_{j,n}^2} \; ; \tag{3.15}$$

the associated position-marginal is the jth marginal of the ring-polymer configurational distribution with density $\exp(-(\beta m_n/2)\, \boldsymbol{q}^{\mathrm{T}}(\Lambda/m + \boldsymbol{\Omega}^2)\boldsymbol{q})$.

For this system, the BAOAB-like update in Sec. 3.2.1 is of the same form as Eq. (3.9) with

$$\begin{aligned} \boldsymbol{\mathcal{M}}_{j,n,\Delta t} &= \boldsymbol{\mathcal{B}}^{1/2}\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2}\boldsymbol{\mathcal{O}}_{j,n,\Delta t}\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2}\boldsymbol{\mathcal{B}}^{1/2} \quad \text{and} \\ \boldsymbol{\mathcal{R}}_{j,n,\Delta t} &= \frac{1-\mathrm{e}^{-4\omega_{j,n}\Delta t}}{\beta m_n}\boldsymbol{\mathcal{B}}^{1/2}\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2}\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}(\boldsymbol{\mathcal{B}}^{1/2}\boldsymbol{\mathcal{S}}_{j,n,\Delta t}^{1/2})^{\mathrm{T}} . \end{aligned} \tag{3.16}$$

As in the case of a constant external potential, the ergodicity of this integrator depends on the spectral radius of $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$. By Theorem 1 in Appendix 3.5 and the fact that

$$\det(\boldsymbol{\mathcal{M}}_{j,n,\Delta t}) = \mathrm{e}^{-2\omega_{j,n}\Delta t} \quad \text{and} \quad \mathrm{Tr}(\boldsymbol{\mathcal{M}}_{j,n,\Delta t}) = \mathcal{A}_{j,n}(1+\mathrm{e}^{-2\omega_{j,n}\Delta t}) ,$$

it follows that condition (**C3**) gives a simple and sufficient condition for ergodicity at any bead number n and hence implies property (**P4**) for the BAOAB-like update specified by Eqs. (3.9) and (3.16). Furthermore, because the matrix $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$ of the corresponding OBABO-like update has equal trace and determinant, condition (**C3**) also guarantees property (**P4**) in that case.[2]

If condition (**C3**) holds, the BAOAB-like update is ergodic with respect to a centered bivariate normal distribution whose covariance matrix $\boldsymbol{\Sigma}_{j,n,\Delta t}$ satisfies the linear equation

$$\boldsymbol{\Sigma}_{j,n,\Delta t} = \boldsymbol{\mathcal{M}}_{j,n,\Delta t}\boldsymbol{\Sigma}_{j,n,\Delta t}\boldsymbol{\mathcal{M}}_{j,n,\Delta t}^{\mathrm{T}} + \boldsymbol{\mathcal{R}}_{j,n,\Delta t} , \tag{3.17}$$

for which the solution is

$$\boldsymbol{\Sigma}_{j,n,\Delta t} = \frac{1}{\beta m_n}\begin{bmatrix} s_{j,n,\Delta t}^2 & 0 \\ 0 & r_{j,n,\Delta t}^2 \end{bmatrix} \tag{3.18}$$

where the variance in the position- and velocity-marginal is respectively $(\beta m_n)^{-1}s_{j,n,\Delta t}^2$ and $(\beta m_n)^{-1}r_{j,n,\Delta t}^2$ with

$$\begin{aligned} s_{j,n,\Delta t}^2 &= \left(\omega_{j,n}^2 + \frac{\Lambda}{m}\frac{\omega_{j,n}\Delta t/2}{\tan\left(\theta(\omega_{j,n}\Delta t)/2\right)}\right)^{-1} \quad \text{and} \\ r_{j,n,\Delta t}^2 &= 1 - \frac{\Delta t^2\Lambda}{4m}\frac{\tan\left(\theta(\omega_{j,n}\Delta t)/2\right)}{\omega_{j,n}\Delta t/2} . \end{aligned} \tag{3.19}$$

Because the tangent function is monotonically increasing on the range of θ specified by condition (**C3**), we have the correspondence

$$0 < s_{j,n,\Delta t}^2 \leq s_{j,n}^2 \quad \text{and} \quad 1 - \frac{\Delta t^2\Lambda}{4m} \leq r_{j,n,\Delta t}^2 < 1 \tag{3.20}$$

---

[2] Condition (**C3**) may be viewed as a relaxation of the sufficient condition for ergodicity given in Eq. (18) of Chapter 2. Indeed, condition (**C3**) implies ergodicity irrespective of the Ornstein–Uhlenbeck friction schedule, whereas Eq. (18) in Chapter 2 does not imply ergodicity for friction schedules that lead to overdamped dynamics.

between the exact and numerical variances of the jth ring-polymer mode. Equation (3.19) reveals that $\theta(x) = 2\arctan(x/2)$ is the unique function that complies with condition (**C3**) and saturates the inequality $s^2_{j,n,\Delta t} \leq s^2_{j,n}$ in Eq. (3.20); consequently, the corresponding BAOAB-like scheme preserves the exact position-marginal in all modes and trivially satisfies property (**P5**). The BCOCB integrator introduced in Chapter 2 corresponds to this choice of θ and thus uniquely provides optimal equilibrium position-marginal accuracy for harmonic external potentials.

To identify other BAOAB-like schemes compliant with condition (**C3**) that satisfy property (**P5**), we examine the overlap between the numerical stationary position-marginal distribution $\mu_{n,\Delta t}$ and the exact distribution μ_n where

$$\mu_n = \prod_{j=1}^{n-1} \mathcal{N}\Big(0, \tfrac{s^2_{j,n}}{\beta m_n}\Big) \quad \text{and} \quad \mu_{n,\Delta t} = \prod_{j=1}^{n-1} \mathcal{N}\Big(0, \tfrac{s^2_{j,n,\Delta t}}{\beta m_n}\Big) \ .$$

Centroid-mode marginals have been suppressed in the definitions of μ_n and $\mu_{n,\Delta t}$. A BAOAB-like scheme is dimension-free if it admits an n-independent upper bound on the distance $d_{\mathrm{TV}}(\mu_n, \mu_{n,\Delta t})$ between μ_n and $\mu_{n,\Delta t}$, where d_{TV} is the total variation metric [40]. In particular, if we require

$$\frac{x}{1+|x|} \leq \theta(x) \leq 2\arctan(x/2) \ \text{ for } \ x > 0 \ , \tag{\textbf{C4}}$$

then we have the dimension-free bound

$$d_{\mathrm{TV}}(\mu_n, \mu_{n,\Delta t}) < \left(\sqrt{\frac{4}{3}\frac{\hbar\beta}{\Delta t}}\right) \frac{\Delta t^2 \Lambda}{m} \ . \tag{3.21}$$

A proof of this claim is provided in Sec. 3.8. Condition (**C4**) ensures that any BAOAB-like integrator with a compliant choice of θ meets property (**P5**).

For OBABO-like schemes, the bound in condition (**C4**) must be tightened to guarantee non-zero overlap between μ_n and $\mu_{n,\Delta t}$ for arbitrarily large n. In particular, replacing $2\arctan(x/2)$ with $\min\{2\arctan(x/2), C\}$ for some $C \in (0,\pi)$ in the upper bound of condition (**C4**) yields a n-independent bound on $d_{\mathrm{TV}}(\mu_n, \mu_{n,\Delta t})$ for all compliant OBABO-like integrators, as can be shown through arguments similar to those in Sec. 3.8.

Taken together, conditions (**C1**)–(**C4**) specify a family of BAOAB-like schemes with dimension-free stability, ergodicity and equilibrium accuracy for applications with harmonic external potentials. Numerical results in Sec. 3.3 suggest that the integrators exhibit similar properties in a more realistic setting with a strongly anharmonic external potential.

3.2.6 Dimension-free convergence to equilibrium of T-RPMD with a harmonic external potential

Beyond ensuring ergodicity of the T-RPMD update in Eq. (3.16), condition (**C3**) leads to explicit dimension-free equilibration rates for compliant schemes. Theorem 4 in Sec. 3.7 proves this result in the infinite-friction limit for ring-polymer modes with arbitrarily high frequency. In detail, the theorem shows that the configurational transition kernel associated with the T-RPMD update of the jth mode in Eq. (3.16) is contractive in the 2-Wasserstein metric [68] and equilibrates any given initial distribution at a rate determined by the function θ, the (external) potential curvature Λ, and the (stable) time-step Δt if condition (**C3**) holds. The rate in Theorem 4, though obtained in the infinite-friction limit, holds for finite friction coefficients $\gamma_{j,n}$ leading to spectral radii $\rho(\boldsymbol{\mathcal{M}}_{j,n,\Delta t}) \le |\mathcal{A}_{j,n}|$, where $\mathcal{A}_{j,n}$ is defined in the display after Eq. (3.13) and $|\mathcal{A}_{j,n}| = \lim_{\gamma_{j,n}\to\infty} \rho(\boldsymbol{\mathcal{M}}_{j,n,\Delta t})$ is the spectral radius at infinite friction.

To illustrate dimension-free convergence, Fig. 3.3 plots the 2-Wasserstein distance between the stationary configurational (i.e., position-marginal) distribution $\mu_{n,\Delta t}$ and the distribution $\mu^k_{n,\Delta t}$ at the kth T-RPMD step evolved from a point mass at the origin using the schemes specified by $\theta(x) = \arccos(\operatorname{sech}(x))$ (Fig. 3.3a), $\theta(x) = 2\arctan(x/2)$ (Fig. 3.3b), and $\theta(x) = \arctan(x)$ (Fig. 3.3c) for a range of bead numbers n. These choices of θ respectively lead to *overdamped*, *critical*, and *Cayley* evolution of the thermostatted free ring polymer under PILE friction (see Sec. 3.2.3), and are identified accordingly in Fig. 3.3. The ring-polymer system considered in Fig. 3.3 approximates the O–H stretch dynamics in liquid water at room temperature with the parameters listed in Sec. 3.3.1. Velocity-marginals were initialized as in the setting of Theorem 4 (see Sec. 3.7), and the position of the jth ring-polymer mode at time $k\Delta t$ follows a centered normal distribution with variance $(\beta m_n)^{-1}(s^k_{j,n,\Delta t})^2$, where

$$(s^k_{j,n,\Delta t})^2 = (\boldsymbol{\mathcal{M}}^k_{j,n,\Delta t})^2_{12} + \beta m_n \sum_{\ell=0}^{k-1} \left(\boldsymbol{\mathcal{M}}^\ell_{j,n,\Delta t} \boldsymbol{\mathcal{R}}_{j,n,\Delta t} (\boldsymbol{\mathcal{M}}^\ell_{j,n,\Delta t})^{\mathrm{T}}\right)_{11} \quad \text{for} \quad k > 0 \, .$$

Here and throughout, the notation $(\,\cdot\,)_{ij}$ chooses the element in the ith row and jth column of the matrix enclosed in the parentheses. The 2-Wasserstein distances in Fig. 3.3 were evaluated using a well-known result for multivariate normal distributions [69].

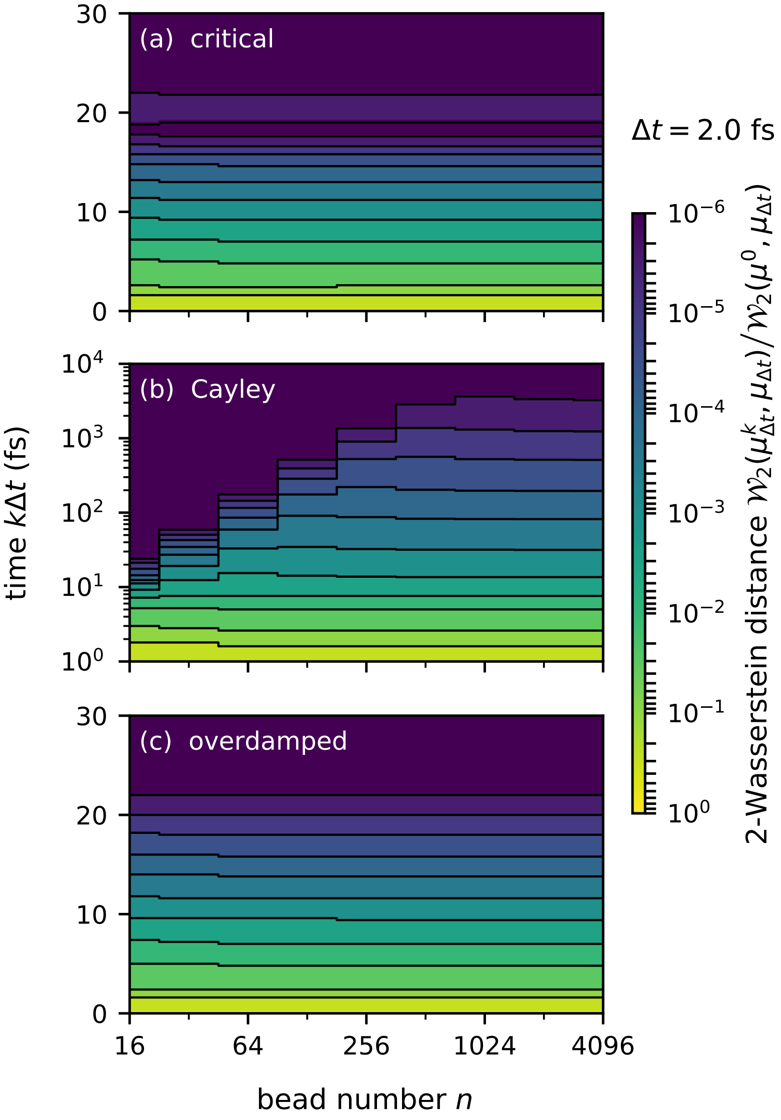

Figure 3.3. Dimension-free convergence to equilibrium of BAOAB-like T-RPMD schemes with a harmonic external potential. The physical parameters of the ring-polymer system (i.e., Λ, m, and β) are listed in Sec. 3.3.1. Panels a, b, and c plot the normalized 2-Wasserstein distance between the configurational ring-polymer distribution at stationarity and at time $k\Delta t$, as evolved via various BAOAB-like schemes from an initial point-mass distribution. Regions with darker color indicate smaller 2-Wasserstein distance to stationarity, and black lines mark iso-distance contours. The contours plateau at some value of n for all tested schemes, which checks that they exhibit dimension-free convergence as predicted by Theorem 4.

Figures 3.3a and 3.3c clearly show that the critical and overdamped schemes converge at dimension-free rates, but this is less evident from Fig. 3.3b for the Cayley scheme. The latter scheme nonetheless displays an n-independent, and hence dimension-free, distance to stationarity at all times $k\Delta t > 0$, indicated by plateauing of the contour lines towards the right of Fig. 3.3b. The ladder-like pattern that precedes this plateau shows a transition from geometric (i.e., fast) to sub-geometric (i.e., slow) convergence upon introducing higher-frequency modes into the ring polymer. The transition manifests with the Cayley scheme because of its aggressive overdamping of the high-frequency modes, which is absent in the other two schemes (see Fig. 3.2).

The example considered in this section illustrates that the equilibration timescale (e.g., the time until the 2-Wasserstein distance decays below 10^{-6}) of the Cayley scheme at large n can dramatically exceed that of other BAOAB-like schemes. Although this negative feature may render the scheme impractical for pathological applications, we find in the next section that the Cayley scheme's superior configurational sampling provides compelling justification for its preferred use in realistic settings.

3.3 Numerical results

The current section provides numerical comparisons of the BAOAB-like T-RPMD integrators in Sec. 3.2, on applications featuring harmonic (Sec. 3.3.1) and anharmonic (Sec. 3.3.2) external potentials. Three choices of θ are considered in the numerical comparisons, namely $\theta(x) = \arctan(x)$, $\theta(x) = \arccos(\operatorname{sech}(x))$, and $\theta(x) = 2\arctan(x/2)$. These choices respectively lead to *overdamped*, *critical*, and *Cayley* evolution of the thermostatted free ring polymer under PILE friction (Sec. 3.2.3), and are identified accordingly throughout this section. It is borne out from the numerical comparisons that the Cayley scheme exhibits superior configurational sampling among the tested schemes in both applications.

3.3.1 One-dimensional quantum harmonic oscillator

In the current section, we numerically integrate Eq. (3.1) with the harmonic potential $V(q) = (\Lambda/2)\, q^2$ using PILE friction (i.e., $\mathbf{\Gamma} = 2\mathbf{\Omega}$), $m = 0.95$ amu, $\sqrt{\Lambda/m} = 3886$ cm^{-1}, and $T = 298$ K. This choice of physical parameters corresponds to a harmonic approximation of the Morse contribution to the O–H bond potential in the q-TIP4P/F force field for water [41], and sets a least upper bound for the T-RPMD stability interval at $\Delta t^{\max} = 2/\sqrt{\Lambda/m} = 2.74$ fs. The simulations reported throughout this section employ the time-step $\Delta t = 0.73 \times \Delta t^{\max} = 2.00$ fs.

Figure 3.4 compares the accuracy and efficiency of various BAOAB-like T-RPMD schemes at equilibrium as a function of the bead number n. For a description of the numerical simulation and statistical estimation procedures used to generate the numerical data (filled circles) in Fig. 3.4, the reader is referred to Sec. 3.11. Figures 3.4a and 3.4c report the mean quantum kinetic energy at equilibrium as per the primitive and virial estimators,

$$\begin{aligned} \mathrm{KE}_n^{\mathrm{pri}}(\boldsymbol{q}) &= \frac{n}{2\beta} - \sum_{j=0}^{n-1} \frac{m_n \omega_n^2}{2} \left(q_{j+1} - q_j\right)^2 \quad \text{and} \\ \mathrm{KE}_n^{\mathrm{vir}}(\boldsymbol{q}) &= \frac{1}{2\beta} + \frac{1}{2} \sum_{j=0}^{n-1} (q_j - \overline{q})\, \partial_{q_j} V_n^{\mathrm{ext}}(\boldsymbol{q}) \ , \end{aligned} \tag{3.22}$$

where $\overline{q} = \frac{1}{n} \sum_{j=0}^{n-1} q_j$ is the centroid position of the n-bead ring polymer. For these two observables, Figs. 3.4b and 3.4d quantify the equilibrium sampling efficiency of the schemes in terms of the integrated autocorrelation time (or normalized asymptotic variance) [70–74]

$$\begin{aligned} \frac{\mathrm{aVar}(\mathsf{O}_n)}{\mathrm{Var}(\mathsf{O}_n)} &= \frac{\lim_{K\to\infty} \mathrm{Var}\Big(\frac{1}{\sqrt{K}} \sum_{k=0}^{K-1} \mathsf{O}_n(\boldsymbol{\xi}^{(k\Delta t)})\Big)}{\mathrm{Var}\big(\mathsf{O}_n\big)} \\ &= 1 + 2 \sum_{k=1}^{\infty} \mathrm{Cor}(\mathsf{O}_n(\boldsymbol{\xi}^{(0)}), \mathsf{O}_n(\boldsymbol{\xi}^{(k\Delta t)})) \ , \end{aligned} \tag{3.23}$$

where O_n is an n-bead observable, $\{\boldsymbol{\xi}^{(k\Delta t)}\}_{k=0}^{\infty} = \{(\boldsymbol{q}^{(k\Delta t)}, \boldsymbol{v}^{(k\Delta t)})\}_{k=0}^{\infty}$ a T-RPMD trajectory initialized at stationarity, $\mathrm{Var}(\mathsf{O}_n)$ the variance of O_n at equilibrium, and $\mathrm{Cor}(\mathsf{O}_n(\boldsymbol{\xi}^{(0)}), \mathsf{O}_n(\boldsymbol{\xi}^{(k\Delta t)}))$ the lag-$k\Delta t$ autocorrelation of O_n along the T-RPMD trajectory. The integrated autocorrelation time of O_n is interpreted as the timescale over which adjacent observations along an equilibrium trajectory become statistically uncorrelated [70–74] and is hence a measure of the efficiency of a T-RPMD scheme at estimating the mean of O_n with respect to the numerically sampled equilibrium distribution. Figures 3.4a-d show that the scheme specified by the Cayley angle (orange) outperforms others in terms of both accuracy and efficiency at estimating the equilibrium average of the quantum kinetic energy observables.

From the perspective of configurational accuracy, the optimality of the Cayley angle displayed in Figs. 3.4a and 3.4c is not surprising in light of the findings in Sec. 3.2.5.

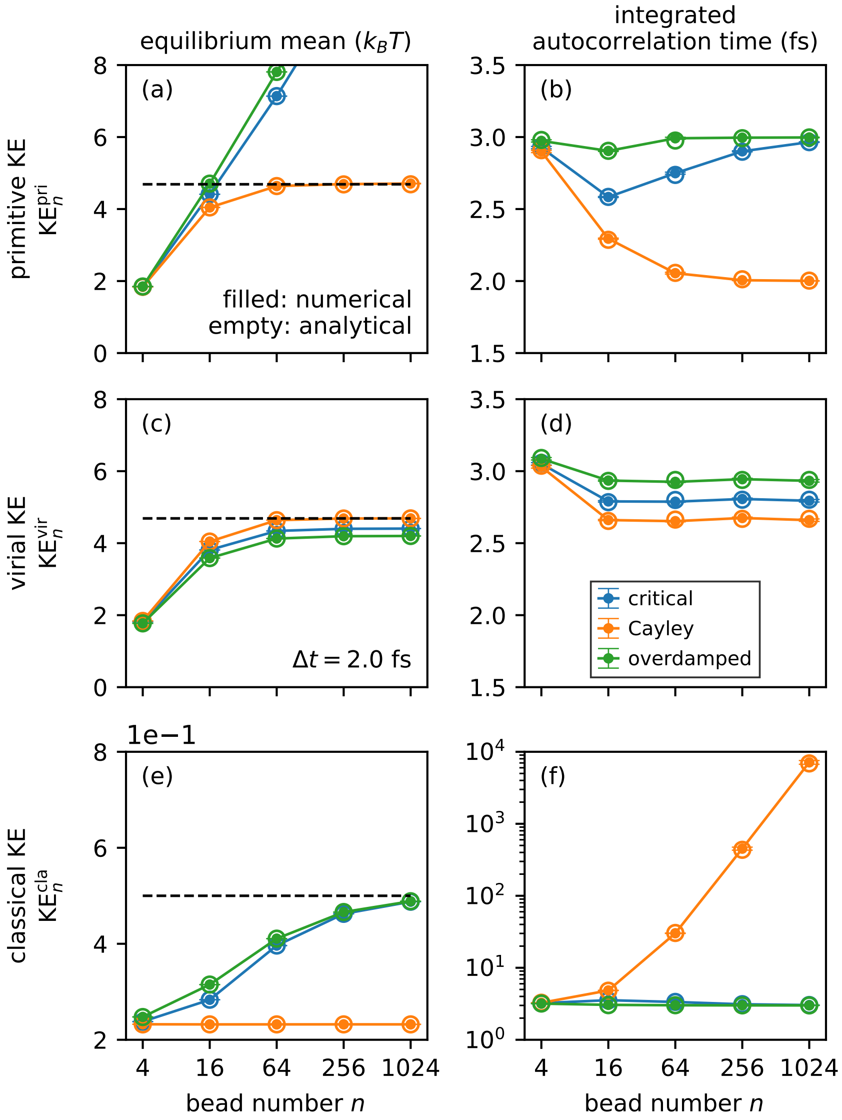

Figure 3.4. Performance at equilibrium of various BAOAB-like T-RPMD schemes applied to the one-dimensional quantum harmonic oscillator with physical parameters listed in Sec. 3.3.1. Panels a, c, and e, respectively, plot the equilibrium mean primitive kinetic energy, virial kinetic energy, and non-centroid classical kinetic energy per mode as a function of bead number n; the corresponding means in the exact infinite-bead limit are plotted as dashed lines. Panels b, d, and f plot the integrated autocorrelation times (Eq. (3.23)) of the respective observables. Exact (resp. numerically estimated) values of the plotted quantities are shown with empty (resp. filled) circles. Numerical estimates were obtained using the protocol described in Sec. 3.11.

Less expected are the results in Figs. 3.4b and 3.4d, which suggest that the Cayley angle is also optimal from the standpoint of configurational sampling efficiency for the quantum kinetic energy observables in Eq. (3.22). Sec. 3.9 supports this conjecture with an analytical result for harmonic external potentials.

Figure 3.4e plots the mean classical kinetic energy at equilibrium as computed from the non-centroid ring-polymer velocities,

$$\mathrm{KE}_n^{\mathrm{cla}}(\boldsymbol{v}) = \frac{m_n}{2(n-1)} \sum_{j=0}^{n-1} (v_j^2 - \overline{v}^2) \approx \frac{1}{2\beta} \, , \tag{3.24}$$

and Fig. 3.4f plots the corresponding integrated autocorrelation time as given by Eq. (3.23). For this observable, the equilibrium accuracy and efficiency of the Cayley scheme are significantly worse than those of the others as n increases. This is a consequence of the strongly overdamped behavior of Cayley T-RPMD at high frequencies (see Fig. 3.2), for which the integrator's ergodicity degrades as its spectral radius approaches unity. Note that this shortcoming of the Cayley scheme presents no adverse implications to the equilibrium sampling of observables that exclusively depend on the ring-polymer configuration, as confirmed by Figs. 3.4a-d.

In summary, Fig. 3.4 establishes that the T-RPMD scheme specified by the Cayley angle provides optimally accurate and efficient configurational sampling *at* equilibrium. To exploit this remarkable feature in practice, the scheme must manifest rapid converge *to* equilibrium when initialized away from it, as is necessary in most realistic applications of T-RPMD. Fortunately, Theorem 4 guarantees that any BAOAB-like scheme compliant with conditions (**C1**)–(**C4**) features a contractive configurational transition kernel for any number of ring-polymer beads, and Fig. 3.3 in Sec. 3.2.6 illustrates this fact for the quantum harmonic oscillator considered in the current section.

3.3.2 Room-temperature liquid water

While theoretical analysis and numerical tests of BAOAB-like T-RPMD schemes in Secs. 3.2.1 and 3.3.1 have focused on harmonic external potentials, the current section demonstrates that the resulting insights carry over to a realistic, strongly anharmonic model of room-temperature liquid water. Our test system is a periodic box containing 32 water molecules at a temperature of 298 K and a density of 0.998 g/cm^3, with potential energy described by the q-TIP4P/F force field [41]. As in Sec. 3.3.1, we compare the performance of various BAOAB-like T-RPMD schemes

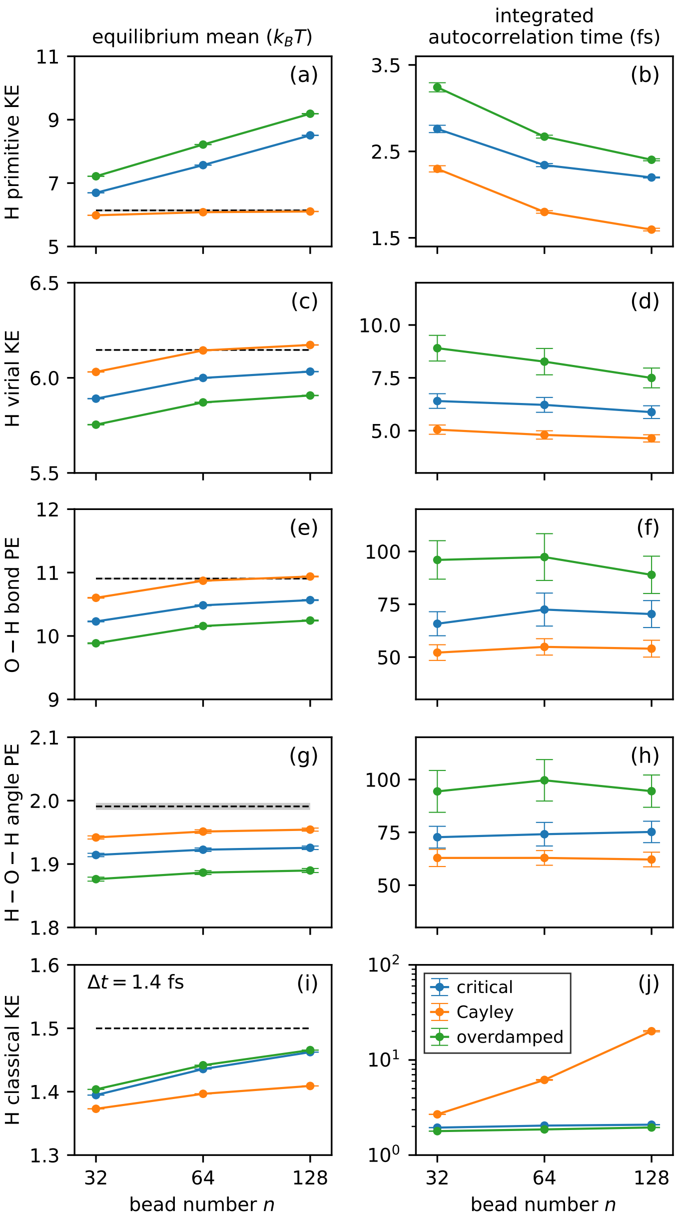

Figure 3.5. Performance of various BAOAB-like T-RPMD schemes applied to q-TIP4P/F liquid water at room temperature. As a function of the bead number n and for a 1.4-fs time-step, panels a and c plot the equilibrium kinetic energy per H atom as per the primitive and virial estimators (Eq. (3.22)), and panels b and d plot the corresponding integrated autocorrelation times. Likewise, panels e and g plot the equilibrium potential energy per H_2O molecule due to the O−H-stretch and H−O−H-bend contributions, as defined in the q-TIP4P/F force field, and the corresponding autocorrelation times are plotted by panels f and h. Finally, panel i plots the classical kinetic energy per H atom computed from the non-centroid velocity estimator (Eq. (3.24)), and panel j plots the corresponding autocorrelation time. The numerical estimates and reference results (dashed lines) were obtained using the protocols described in Sec. 3.11.

for integrating the many-dimensional analogue of Eq. (3.1) with PILE friction, using the simulation time-step $\Delta t = 1.4$ fs in all simulations. Numerical tests reported in Sec. 3.10 show that this value of Δt closely approximates the upper limit of the Verlet (i.e., $n = 1$) stability interval for q-TIP4P/F liquid water. In agreement with Sec. 3.3.1, the experiments reveal that among the tested T-RPMD schemes, the Cayley scheme offers superior configurational sampling. For details on the numerical simulation and statistical estimation procedures used to generate the data presented in this section, the reader is referred to Sec. 3.11.

Figure 3.5 compares the equilibrium accuracy achieved by the tested schemes in terms of the quantum and classical kinetic energy per hydrogen atom (Figs. 3.5a, 3.5c, and 3.5i) and the intramolecular potential energy per water molecule (Figs. 3.5e and 3.5g); also plotted are the respective integrated autocorrelation times as a function of bead number n. The kinetic energy estimates in Figs. 3.5a and 3.5c exhibit similar trends to those seen in Fig. 3.4 for the one-dimensional harmonic oscillator. In particular, the T-RPMD scheme specified by the Cayley angle outperforms others in terms of quantum kinetic energy accuracy as n increases, most outstandingly with a highly accurate primitive kinetic energy estimate despite the large time-step employed. Still in close agreement with the harmonic oscillator results, Figs. 3.5b and 3.5d show that the Cayley scheme displays the shortest integrated autocorrelation time among the tested schemes for the quantum kinetic energy observables. Similar trends manifest in the intramolecular potential energy averages and their autocorrelation times (Figs. 3.5e-h), where the Cayley scheme also achieves superior accuracy and efficiency. Finally, Figs. 3.5i and 3.5j confirm that the relative performance of the compared schemes in terms of velocity-marginal sampling is qualitatively consistent with the harmonic results. Taken together, the results in Fig. 3.5 suggest that the superiority of the Cayley scheme for configurational sampling, proven in the model setting of a harmonic external potential, is also reflected in realistic applications.

In a final numerical test, Fig. 3.6 confirms that the sampling advantages of the Cayley T-RPMD scheme are obtained without downside in the estimation of dynamical quantities of typical interest. Specifically, Fig. 3.6b shows (unnormalized) infrared absorption spectra for room-temperature liquid water, computed from the 128-bead T-RPMD trajectories used to generate Fig. 3.5 using linear response theory and the T-RPMD approximation to real-time quantum dynamics [20, 42]. Linear response dictates that the absorption spectrum is proportional to $\omega^2\tilde{\mathcal{I}}(\omega)$, where

$\tilde{\mathcal{I}}(\omega) = \int_{\mathbb{R}} \mathrm{d}t\, \mathrm{e}^{-i\omega t} \tilde{C}_{\mu\mu}(t)$ is the Fourier transform of the quantum mechanical Kubo-transformed dipole autocorrelation function $\tilde{C}_{\mu\mu}(t)$. The latter is approximated within the T-RPMD framework [12, 16] by $\tilde{C}_{\mu\mu}(t) \approx \frac{1}{N_{\mathrm{H_2O}}} \sum_{i=1}^{N_{\mathrm{H_2O}}} \langle \overline{\mu}_i(t) \cdot \overline{\mu}_i(0) \rangle$, where $N_{\mathrm{H_2O}}$ is the number of molecules in the liquid, $\overline{\mu}_i(t)$ is the bead-averaged dipole moment vector of molecule i at time t, and the covariance $\langle \overline{\mu}_i(t) \cdot \overline{\mu}_i(0) \rangle$ is estimated from a stationary T-RPMD trajectory as indicated in Sec. 3.11. Figure 3.6a plots the T-RPMD estimates of $\tilde{C}_{\mu\mu}(t)$ leading to the absorption spectra in Fig. 3.6b. On the scale in which the absorption spectrum exhibits its key features, the spectra in Fig. 3.6b show very minor qualitative discrepancies. A similar conclusion holds for Fig. 3.6c, where the T-RPMD approximation of the Kubo-transformed velocity autocovariance function $\tilde{C}_{vv}(t) \approx \frac{1}{N_{\mathrm{H_2O}}} \sum_{i=1}^{N_{\mathrm{H_2O}}} \langle \overline{v}_i(t) \cdot \overline{v}_i(0) \rangle$ is plotted for the three tested T-RPMD schemes. Collectively, these observations indicate that the accuracy of dynamical properties computed with BAOAB-like schemes is not significantly affected by the particular θ employed if conditions (**C1**)–(**C4**) in Sec. 3.2 are met. This result is expected due to the fact that the considered dynamical properties depend on bead-averaged (i.e., centroid-mode) coordinates, whose evolution is largely independent of the choice of θ under weak coupling between the centroid and non-centroid ring-polymer modes.

3.4 Summary

Recent works showed that strong stability [26] and dimensionality [75] are essential features of a robust T-RPMD integration scheme that many widespread integrators do not possess. A T-RPMD scheme with these features, denoted BCOCB, was introduced via a simple and inexpensive Cayley modification of the free ring-polymer update (i.e., the “A” sub-step) of the standard BAOAB integrator. The BCOCB scheme was then shown to dramatically outperform BAOAB at estimating static and dynamic properties of various systems with remarkable accuracy at unprecedented time-steps.

The current chapter generalizes beyond the Cayley modification by introducing a simple parameterization of the free ring-polymer update and a corresponding family of strongly stable and dimension-free modifications of the BAOAB scheme. Among these schemes lies BCOCB, which is found to exhibit superior configurational sampling despite exhibiting worse accuracy and efficiency for observables that depend on the non-centroid ring-polymer velocities. This conclusion is obtained theoretically via exhaustive analysis of a harmonic model, and numerically via simulation

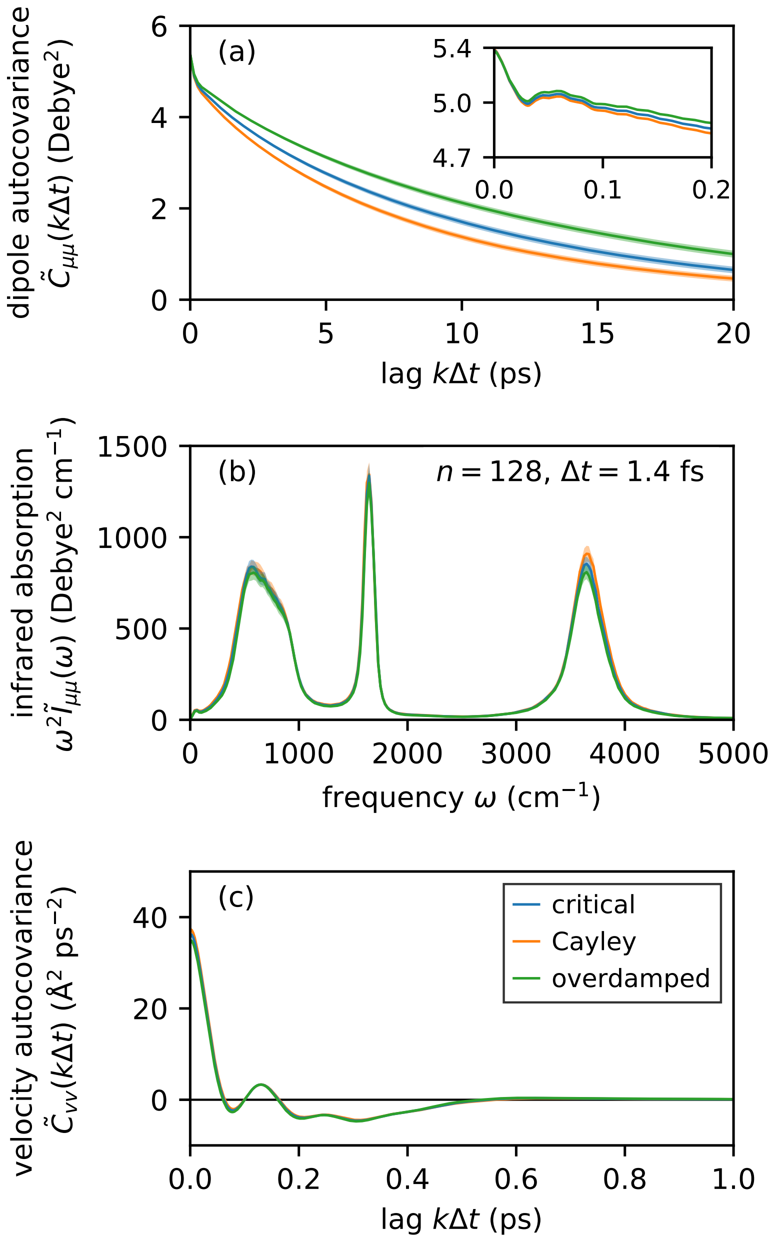

Figure 3.6. Molecular dipole autocovariance function (panel a), corresponding infrared absorption spectrum (panel b), and molecular velocity autocovariance function (panel c) in room-temperature liquid water for various BAOAB-like T-RPMD schemes. The plotted quantities autocovariance exhibit minor qualitative discrepancies across schemes, which suggests that all schemes compliant with conditions (**C1**)–(**C4**) exhibit comparable accuracy in the computation of dynamical properties. Numerical estimates of the autocovariance functions were obtained using the protocol described in Sec. 3.11.

of a realistic quantum-mechanical model of liquid water at room temperature. In this way, the chapter supports the outstanding utility of BCOCB for accurate and efficient equilibrium simulation of condensed-phase systems with T-RPMD.

To conclude, we stress that implementing BCOCB or any of the new dimension-free and strongly-stable schemes adds no cost, no algorithmic parameters and insignificant coding overhead relative to the standard BAOAB integrator. The modified integrators thus provide "turnkey" means to significantly improve the accuracy and stability of existing RPMD and T-RPMD implementations [76, 77].

Acknowledgments. J. L. R.-R. and J. S. contributed equally to the work presented in this chapter. The work was supported in part by the U.S. Department of Energy (DE-SC0019390) and the National Institutes of Health (R01GM125063) N. B.-R. acknowledges support by the Alexander von Humboldt foundation and the National Science Foundation (DMS-1816378).

3.5 Appendix A: Necessary and sufficient condition for eigenvalues of a 2×2 real matrix to be inside the unit circle

This section provides a proof of the standard result used in Secs. 3.2.3 and 3.2.5 to infer ergodicity of the T-RPMD update for free and harmonically-confined ring polymers.

Theorem 1. *The spectral radius of a 2×2 real matrix $\boldsymbol{M}$ is strictly less than one if and only if*

$$|\operatorname{Tr}(\boldsymbol{M})| < 1 + \det(\boldsymbol{M}) < 2\ . \tag{3.25}$$

Figure 3.7 plots eigenvalue pairs λ_1, λ_2 that satisfy Eq. (3.25) for a fixed value of $\det(\boldsymbol{M}) = \lambda_1\lambda_2$. Note that the spectral radius of $\boldsymbol{M}$ is minimized when λ_1 and λ_2 are on the circle with radius $r = \sqrt{\det(\boldsymbol{M})}$.

Proof. Let λ_1, λ_2 be the (possibly complex) eigenvalues of $\boldsymbol{M}$. By definition, the spectral radius of $\boldsymbol{M}$ is $\max(|\lambda_1|, |\lambda_2|) =: \rho$. Since $\boldsymbol{M}$ is real, both $\operatorname{Tr}(\boldsymbol{M}) = \lambda_1 + \lambda_2$ and $\det(\boldsymbol{M}) = \lambda_1\lambda_2$ are real. Thus, either:

1. λ_1, λ_2 are a complex conjugate pair; or,
2. λ_1, λ_2 are both real.

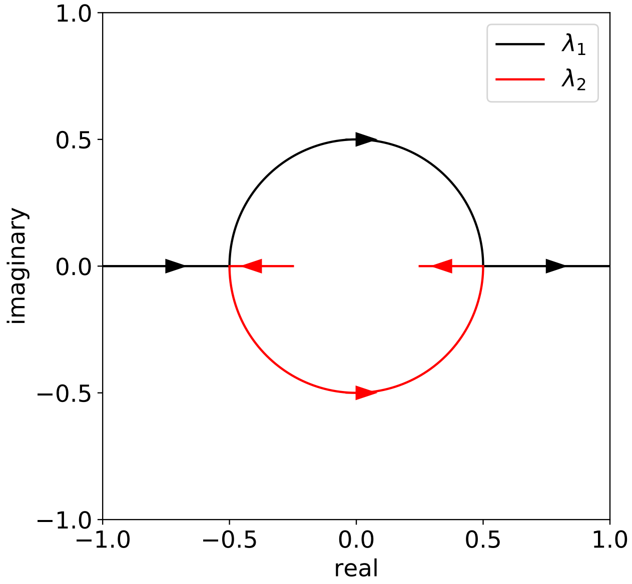

Figure 3.7. All possible eigenvalue pairs λ_1, λ_2 of a matrix $\boldsymbol{M}$ that satisfies Eq. (3.25) with $\det(\boldsymbol{M}) = \lambda_1\lambda_2 = 1/4$. The eigenvalue pairs either lie on the circle with radius $r = 1/2$ or are both real, and in the former case, the spectral radius of $\boldsymbol{M}$ is minimal.

In the first case, $\lambda_1 = a + ib$ and $\lambda_2 = a - ib$ for some real numbers a and b with $b \neq 0$, and hence, $\det(\boldsymbol{M}) = \lambda_1\lambda_2 = a^2 + b^2 > 0$, and $\rho = |\lambda_1| = |\lambda_2| = \sqrt{a^2 + b^2}$, i.e., the eigenvalues lie on the circle with radius $\rho = \sqrt{a^2 + b^2} = \sqrt{\det(\boldsymbol{M})}$. In this case, the first inequality in Eq. (3.25) holds since $b \neq 0$ implies

$$|\operatorname{Tr}(\boldsymbol{M})| = 2|a| < 2\rho \leq 1 + \rho^2 = 1 + \det(\boldsymbol{M}) \ .$$

Hence, Eq. (3.25) is equivalent to $1 + \det(\boldsymbol{M}) < 2$ or $\rho < 1$.

In the second case, λ_1, λ_2 are both real, and the condition $|\operatorname{Tr}(\boldsymbol{M})| < 1 + \det(\boldsymbol{M})$ is equivalent to

$$\begin{aligned} 1 + \lambda_1\lambda_2 + \lambda_1 + \lambda_2 &= (1 + \lambda_1)(1 + \lambda_2) > 0 \ , \text{ and} \\ 1 + \lambda_1\lambda_2 - \lambda_1 - \lambda_2 &= (1 - \lambda_1)(1 - \lambda_2) > 0 \ . \end{aligned}$$

Together with $\det(\boldsymbol{M}) = \lambda_1\lambda_2 < 1$, these conditions imply $\rho = \max(|\lambda_1|, |\lambda_2|) < 1$. $\square$

3.6 Appendix B: Stability condition for harmonic external potentials

This section proves that condition (**C3**) implies property (**P3**), as claimed in Sec. 3.2.4. For notational brevity, we define

$$A(x) := \cos(\theta(x)) - \frac{\Delta t^2(\Lambda/m)}{2}\frac{\sin(\theta(x))}{x} \ .$$

Note that $A(x)$ is equal to $\mathcal{A}_{j,n}$ in the display under Eq. (3.13) if $x = \omega_{j,n}\Delta t$.

Theorem 2. *For any* $\alpha^\star > 0$, (**F2**) *implies* (**F1**).

- (**F1**) *For all* $\Lambda \geq 0$, $m > 0$ *and* $\Delta t > 0$ *satisfying* $\Delta t^2 \Lambda / m < \alpha^\star$, *the function* θ *satisfies*

$$|A(x)| < 1 \ \textit{for} \ \ x > 0 \ .$$

- (**F2**) *The function* θ *satisfies:*

$$0 < \theta(x) < 2\arctan(2x/\alpha^\star) \ \textit{for} \ \ x > 0 \ .$$

Proof. Let $\alpha = \Delta t^2(\Lambda/m)$. For notational brevity, define

$$\phi_\alpha(x) := \arctan(\alpha/(2x)) \ \ \text{for} \ \ x > 0 \ .$$

By the harmonic addition identity

$$\cos(\theta) - \tan(\phi_\alpha)\sin(\theta) = \frac{\cos(\theta + \phi_\alpha)}{\cos(\phi_\alpha)} \ ,$$

note that (**F1**) can be rewritten as

$$\left| \frac{\cos(\theta(x) + \phi_\alpha(x))}{\cos(\phi_\alpha(x))} \right| < 1 \ \ \text{for} \ \ x > 0 \ , \ 0 < \alpha < \alpha^\star \ . \tag{3.26}$$

For $0 < \theta(x) < \pi$, Eq. (3.26) holds if and only if

$$\phi_\alpha(x) < \theta(x) + \phi_\alpha(x) < \pi - \phi_\alpha(x) \ ,$$

which can be rewritten as

$$0 < \theta(x) < 2\arctan(2x/\alpha) \ , \tag{3.27}$$

where we used the identity

$$\pi - 2\arctan(x) = 2\arctan(1/x) \ \ \text{valid for} \ \ x > 0 \ .$$

Since $\arctan$ is monotone increasing, and $0 < \alpha < \alpha^\star$ by assumption, we may conclude that

$$0 < \theta(x) < 2\arctan(2x/\alpha^\star) < 2\arctan(2x/\alpha)\ .$$

Thus, if (**F1**) holds, then Eq. (3.27) holds and therefore (**F1**) holds. $\square$

Fix $\epsilon \in (0,1)$. Since Theorem 2 is true for arbitrary $\alpha^\star$, if we take $\alpha^\star = 4-\epsilon$, then the theorem holds with $\Delta t^2\Lambda/m < 4-\epsilon$ in Theorem 2 (**F1**), and $\theta(x) < 2\arctan(2x/(4-\epsilon))$ in Theorem 2 (**F1**). Since $\epsilon > 0$ is arbitrary, and $\arctan$ is monotone increasing, we can conclude that the theorem holds with $\Delta t^2\Lambda/m < 4$ and $\theta(x) \leq 2\arctan(x/2)$. Summarizing,

Corollary 3. *Suppose that the function θ satisfies*

$$0 < \theta(x) \leq 2\arctan(x/2) \ \ \textit{for} \ \ x > 0\ .$$

Then for all $\Lambda \geq 0$, $m > 0$ and $\Delta t > 0$ satisfying $\Delta t^2\Lambda/m < 4$, we have

$$|A(x)| \ < \ 1 \ \ \textit{for} \ \ x > 0\ .$$

3.7 Appendix C: Dimension-free quantitative contraction rate for harmonic external potentials in the infinite-friction limit

In the infinite-friction limit, Eq. (3.16) simplifies to

$$\begin{aligned}\boldsymbol{\mathcal{M}}_{j,n,\Delta t} &= \boldsymbol{\mathcal{B}}^{1/2}\boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t}\begin{bmatrix}1 & 0\\ 0 & 0\end{bmatrix}\boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t}\boldsymbol{\mathcal{B}}^{1/2} \quad \text{and} \\ \boldsymbol{\mathcal{R}}_{j,n,\Delta t} &= \frac{1}{\beta m_n}\boldsymbol{\mathcal{B}}^{1/2}\boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t}\begin{bmatrix}0 & 0\\ 0 & 1\end{bmatrix}(\boldsymbol{\mathcal{B}}^{1/2}\boldsymbol{\mathcal{S}}^{1/2}_{j,n,\Delta t})^{\mathrm{T}}\ .\end{aligned}$$

The kth step of the corresponding T-RPMD integrator can be written compactly as

$$\begin{bmatrix}\varrho_j^{(k)}\\ \varphi_j^{(k)}\end{bmatrix} = \boldsymbol{\mathcal{M}}_{j,n,\Delta t}\begin{bmatrix}\varrho_j^{(k-1)}\\ \varphi_j^{(k-1)}\end{bmatrix} + \boldsymbol{\mathcal{R}}^{1/2}_{j,n,\Delta t}\begin{bmatrix}\xi_j^{(k-1)}\\ \eta_j^{(k-1)}\end{bmatrix}\ ,$$

where $\xi_j^{(k-1)}$ and $\eta_j^{(k-1)}$ are independent standard normal random variables. Suppose that the initial velocity is drawn from the Maxwell–Boltzmann distribution, $\varphi_j^{(0)} \sim \mathcal{N}(0,(\beta m_n)^{-1})$, and the initial position is drawn from an arbitrary distribution μ_j on $\mathbb{R}$, $\varrho_j^{(0)} \sim \mu_j$. Let $p^k_{j,n}$ denote the k-step transition kernel of the position-marginal, i.e., $\mu_j p^k_{j,n}$ is the probability distribution of $\varrho_j^{(k)}$ with $\varrho_j^{(0)} \sim \mu_j$.

The next theorem shows that starting from any two initial distributions μ_j and ν_j on $\mathbb{R}$, the distance between the distributions $\mu_j p_{j,n}^k$ and $\nu_j p_{j,n}^k$ is contractive. We quantify the distance between these distributions in terms of the 2-Wasserstein metric. For two probability distributions μ and ν on $\mathbb{R}$, the 2-Wasserstein distance between μ and ν is defined as:

$$\mathcal{W}_2(\mu,\nu) := \Big(\inf_{\substack{X\sim\mu \\ Y\sim\nu}} \mathbb{E}(|X-Y|^2) \Big)^{1/2} ,$$

where the infimum is taken over all bivariate random variables (X, Y) such that $X \sim \mu$ and $Y \sim \nu$ [68], and where $\mathbb{E}(\cdot)$ denotes a mathematical expectation.

Theorem 4. *Suppose that the function θ satisfies*

$$0 < \theta(x) \le 2\arctan(x/2) \ \textit{ for } \ x > 0 \ .$$

Then for all $k > 1$, $\Lambda \ge 0$, $m > 0$ and $\Delta t > 0$ satisfying $\Delta t^2\Lambda/m < 4$, and for all initial distributions μ_j and ν_j on $\mathbb{R}$,

$$\mathcal{W}_2(\mu_j p_{j,n}^k, \nu_j p_{j,n}^k) \ \le \ \begin{cases} A(\omega_{j,n}\Delta t)^{k-1}\mathcal{W}_2(\mu_j,\nu_j) & \textit{if } A(\omega_{j,n}\Delta t) > 0, \\ \dfrac{1}{2(k-1)}\mathcal{W}_2(\mu_j,\nu_j) & \textit{else.} \end{cases} \tag{3.28}$$

Proof. In the infinite-friction limit, the eigenvalues of $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$ are $\{0, A(\omega_{j,n}\Delta t)\}$, where $A(x)$ is defined in Appendix 3.6. Let $\varrho_j^{(0)} \sim \mu_j$ and $\tilde{\varrho}_j^{(0)} \sim \nu_j$ be an optimal coupling of μ_j and ν_j, which implies that $\mathcal{W}_2(\mu_j,\nu_j) = \mathbb{E}(|\varrho_j^{(0)} - \tilde{\varrho}_j^{(0)}|^2)^{1/2}$. Conditional on $\varrho_j^{(0)}$ and $\tilde{\varrho}_j^{(0)}$, $\varrho_j^{(k)}$ and $\tilde{\varrho}_j^{(k)}$ are Gaussian random variables with equal variances, but different means. By a well-known result for the 2-Wasserstein distance between Gaussian distributions [69],

$$\begin{aligned} \mathcal{W}_2(\mu_j p_{j,n}^k, \nu_j p_{j,n}^k)^2 \ &= \ |A(\omega_{j,n}\Delta t)|^{2(k-1)} (\boldsymbol{\mathcal{M}}_{j,n,\Delta t})_{11}^2 \mathcal{W}_2(\mu_j,\nu_j)^2 \\ &= \ |A(\omega_{j,n}\Delta t)|^{2(k-1)} \frac{(1 + A(\omega_{j,n}\Delta t))^2}{4} \mathcal{W}_2(\mu_j,\nu_j)^2 \ , \end{aligned} \tag{3.29}$$

where we used $(\boldsymbol{\mathcal{M}}_{j,n,\Delta t})_{11} = (1 + A(\omega_{j,n}\Delta t))/2$.

Now we distinguish between two cases. In the case where $A(\omega_{j,n}\Delta t) > 0$, we obtain the required result since $|A(\omega_{j,n}\Delta t)| < 1$ by Corollary 3, and therefore,

$$\frac{(1 + A(\omega_{j,n}\Delta t))^2}{4} \le 1 \ . \tag{3.30}$$

Otherwise, for $-1 < A(\omega_{j,n}\Delta t) \leq 0$ the quantity $|A(\omega_{j,n}\Delta t)|^{2(k-1)}(1+A(\omega_{j,n}\Delta t))^2$ is maximized at $\big((1-k)/k\big)^{2k}(k-1)^{-2}$, and therefore,

$$|A(\omega_{j,n}\Delta t)|^{2(k-1)} \frac{(1+A(\omega_{j,n}\Delta t))^2}{4} \leq \frac{1}{4(k-1)^2} \,. \tag{3.31}$$

Inserting Eq. (3.30) and Eq. (3.31) into Eq. (3.29), and then taking square roots, gives the required result. □

3.8 Appendix D: Total variation bound on the equilibrium accuracy error for harmonic external potentials

In this section, we show that Eq. (3.21) follows from conditions (**C1**)–(**C4**) in the setting of Sec. 3.2.5. It is helpful to recall the quantities

$$\omega_j = \lim_{n\to\infty} \omega_{j,n} = \begin{cases} \dfrac{\pi j}{\hbar\beta} & \text{if } j \text{ is even} \,, \\ \dfrac{\pi(j+1)}{\hbar\beta} & \text{else} \,. \end{cases} \tag{3.32}$$

In the following, $\mu_{j,n,\Delta t}$ and $\mu_{j,n}$ respectively denote the jth factor of the product distributions $\mu_{n,\Delta t}$ and μ_n introduced in Sec. 3.2.5.

Theorem 5. *Suppose that the function θ satisfies conditions* (**C1**)–(**C4**). *Then for all $\Lambda \geq 0$, $m > 0$ and $\Delta t > 0$ satisfying $\Delta t^2 \Lambda/m < 4$, the total variation distance between μ_n and $\mu_{n,\Delta t}$ is bounded as in Eq.* (3.21).

Proof. Subadditivity of the total variation distance d_{TV} between product distributions and its equivalence with the Hellinger distance [40] d_{H} lead to the inequalities

$$\begin{aligned} d_{\mathrm{TV}}(\mu_n, \mu_{n,\Delta t})^2 &\leq \sum_{j=1}^{n-1} d_{\mathrm{TV}}(\mu_{j,n}, \mu_{j,n,\Delta t})^2 \\ &\leq \sum_{j=1}^{n-1} 2\, d_{\mathrm{H}}(\mu_{j,n}, \mu_{j,n,\Delta t})^2 \leq \sum_{j=1}^{n-1} \frac{2(s_{j,n} - s_{j,n,\Delta t})^2}{(s_{j,n}^2 + s_{j,n,\Delta t}^2)} \\ &\leq \sum_{j=1}^{n-1} \left(1 - \frac{s_{j,n}}{s_{j,n,\Delta t}}\right)^2 \leq \sum_{j=1}^{n-1} \left(1 - \frac{s_{j,n}^2}{s_{j,n,\Delta t}^2}\right)^2 \,, \end{aligned} \tag{3.33}$$

where the second-to-last step uses Eq. (3.20) and the last step uses the elementary inequality $(1-x^2)^2 \geq (1-x)^2$ valid for all $x \geq 0$.

Since $\tan$ increases superlinearly on the interval $(0, \pi)$, we have

$$\theta(x)/2 \leq \tan(\theta(x)/2) \leq x/2 \ \text{ for } \ x > 0 \,,$$

where the second inequality uses (**C3**). Consequently, the jth summand in Eq. (3.33) admits the bound

$$\begin{aligned}\left(1-\frac{s_{j,n}^2}{s_{j,n,\Delta t}^2}\right)^2 &= \left(\frac{\Lambda/m}{\omega_{j,n}^2+\Lambda/m}\left(\frac{\omega_{j,n}\Delta t/2}{\tan\left(\theta(\omega_{j,n}\Delta t)/2\right)}-1\right)\right)^2\\ &\le \left(\frac{\Delta t^2\Lambda/m}{(\omega_{j,n}\Delta t)^2}\left(\frac{\omega_{j,n}\Delta t}{\theta(\omega_{j,n}\Delta t)}-1\right)\right)^2\\ &\le \left(\frac{\Delta t^2\Lambda}{m}\right)^2\frac{1}{(\omega_{j,n}\Delta t)^2}\ ,\end{aligned}$$

where the last line uses the lower bound in (**C4**). Using that for any even positive integer n

$$\sum_{j=1}^{n-1}\frac{1}{\omega_{j,n}^2}<\lim_{n\to\infty}\sum_{j=1}^{n-1}\frac{1}{\omega_{j,n}^2}=\sum_{j=1}^{\infty}\frac{1}{\omega_j^2}<\left(\frac{\hbar\beta}{\pi}\right)^2\sum_{j=1}^{\infty}\frac{2}{j^2}\ ,$$

where we used the definition in Eq. (3.32), the bound in Eq. (3.33) becomes

$$d_{\mathrm{TV}}(\mu_n,\mu_{n,\Delta t})^2<\left(\frac{\Delta t^2\Lambda}{m}\right)^2\left(\frac{\hbar\beta}{\pi\Delta t}\right)^2\sum_{j=1}^{\infty}\frac{2}{j^2}\ .$$

Taking square roots and using the Riemann zeta function [78] to evaluate the infinite sum yields Eq. (3.21). □

3.9 Appendix E: Asymptotic variance of kinetic energy observables for harmonic external potentials in the infinite-friction limit

In Sec. 3.3.1, Figs. 3.4b and 3.4d show that the T-RPMD scheme specified by $\theta(x)=2\arctan(x/2)$, which coincides with the Cayley-modified BAOAB scheme introduced in Chapter 2, provides the smallest integrated autocorrelation time (Eq. (3.23)) for quantum kinetic energy observables (Eq. (3.22)) among several schemes with properties (**P1**)–(**P5**). In this section, we show that this scheme minimizes an upper bound (Eq. (3.35)) on the integrated autocorrelation time of the quantum kinetic energy among all dimension-free and strongly-stable BAOAB-like schemes for harmonic external potentials.

To this end, note that for a n-bead thermostatted ring polymer with external potential

$V_n^{\text{ext}}(\boldsymbol{q}) = \frac{\Lambda}{2n}|\boldsymbol{q}|^2$, Eq. (3.22) can be rewritten

$$\begin{aligned} \mathrm{KE}_n^{\text{pri}}(\boldsymbol{\varrho}) &= \frac{n}{2\beta} - \sum_{j=1}^{n-1} \frac{m_n \omega_{j,n}^2}{2} \varrho_j^2 \quad \text{and} \\ \mathrm{KE}_n^{\text{vir}}(\boldsymbol{\varrho}) &= \frac{1}{2\beta} + \sum_{j=1}^{n-1} \frac{\Lambda}{2n} \varrho_j^2 \end{aligned} \tag{3.34}$$

where $\boldsymbol{\varrho}$ is defined in Eq. (3.5). In the following, we denote both observables in Eq. (3.34) as KE_n and distinguish between the two as needed.

To control the integrated autocorrelation time of KE_n, we need the stationary autocorrelation $\mathrm{Cor}\big(\mathrm{KE}_n(\boldsymbol{\varrho}^{(0)}), \mathrm{KE}_n(\boldsymbol{\varrho}^{(k\Delta t)})\big)$ for $k \geq 0$. Note that the distributions of $\boldsymbol{\varrho}^{(k\Delta t)}$ and $\boldsymbol{\varrho}^{(0)}$ are equal by stationarity, and that components $(\varrho_j)_{j=0}^{n-1}$ are uncorrelated in a harmonic external potential. Thus,

$$\mathrm{Cor}\big(\mathrm{KE}_n(\boldsymbol{\varrho}^{(0)}),\, \mathrm{KE}_n(\boldsymbol{\varrho}^{(k\Delta t)})\big) = \sum_{j=1}^{n-1} w_{j,n} \mathrm{Cor}\big(|\varrho_j^{(0)}|^2, |\varrho_j^{(k\Delta t)}|^2\big) \ ,$$

where

$$w_{j,n} = \frac{\kappa_{j,n}^2 \mathrm{Var}\big(|\varrho_j^{(0)}|^2\big)}{\sum_{i=1}^{n-1} \kappa_{i,n}^2 \mathrm{Var}\big(|\varrho_i^{(0)}|^2\big)}$$

and

$$\kappa_{j,n} = \begin{cases} \dfrac{m_n \omega_{j,n}^2}{2} & \text{for } \mathrm{KE}_n^{\text{pri}} \ , \\ \dfrac{\Lambda}{2n} & \text{for } \mathrm{KE}_n^{\text{vir}} \ . \end{cases}$$

If the evolution of the ring polymer is governed by the BAOAB-like update in Eq. (3.9), then the jth mode satisfies

$$\begin{aligned} \mathrm{Cor}\big(|\varrho_j^{(0)}|^2, |\varrho_j^{(k\Delta t)}|^2\big) &= \frac{\mathrm{Cov}\big(|\varrho_j^{(0)}|^2, |\varrho_j^{(k\Delta t)}|^2\big)}{\mathrm{Var}\big(|\varrho_j^{(0)}|^2\big)} \\ &= (\boldsymbol{\mathcal{M}}_{j,n,\Delta t}^k)_{11}^2 \ , \end{aligned}$$

where we used that the phase $[\varrho_j^{(k\Delta t)}, \varphi_j^{(k\Delta t)}]^{\mathrm{T}}$ follows a centered Gaussian distribution with covariance given in Eq. (3.18) for all $k \geq 0$. Therefore, in the infinite-friction limit where $\boldsymbol{\mathcal{M}}_{j,n,\Delta t}$ is given in Appendix 3.7, the integrated autocorrelation time of KE_n evaluates to

$$\begin{aligned} \frac{\mathrm{aVar}(\mathrm{KE}_n)}{\mathrm{Var}(\mathrm{KE}_n)} &= 1 + 2 \sum_{j=1}^{n-1} w_{j,n} \sum_{k=1}^{\infty} (\boldsymbol{\mathcal{M}}_{j,n,\Delta t}^k)_{11}^2 \\ &\leq 1 + \frac{1}{2} \max_{1 \leq j \leq n-1} \left| \frac{1 + A(\omega_{j,n}\Delta t)}{1 - A(\omega_{j,n}\Delta t)} \right| \ , \end{aligned} \tag{3.35}$$

where simplification of $(\boldsymbol{\mathcal{M}}^{k}_{j,n,\Delta t})_{11}$ was aided by the Cayley–Hamilton theorem for 2×2 matrices [79], $A(x)$ is defined in Appendix 3.6, and in the last line we used that $\sum_{j=1}^{n-1} w_{j,n} = 1$. Equation (3.35) states that the integrated autocorrelation time of KE_n can only be as small as that of the component $|\varrho_j|^2$ exhibiting the slowest decorrelation at stationarity.

Having derived Eq. (3.35), we now prove our claim for this section. Let $x := \omega_{j,n}\Delta t > 0$ and $\alpha := \Delta t^2 \Lambda/m \in (0,4)$. For fixed x and α, the function $A(x) = \cos(\theta(x)) - \frac{\alpha}{2x}\sin(\theta(x))$ monotonically decreases toward -1 as the angle $\theta(x)$ increases toward π. Consequently, the function $\big|\big(1 + A(x)\big)/\big(1 - A(x)\big)\big|$ decreases (toward 0) as $\theta(x)$ increases (toward π), but condition (**C3**) requires that $\theta(x) \leq 2\arctan(x/2)$ to achieve stable evolution. Therefore, because it yields the largest stable angle, the choice $\theta(x) = 2\arctan(x/2)$, corresponding to the Cayley angle, minimizes the upper bound in Eq. (3.35).

A similar argument can be made to support the conjecture, suggested by Fig. 3.4f, that the non-centroid velocity estimator for the classical kinetic energy $\mathrm{KE}_n^{\mathrm{cla}}$ in Eq. (3.24), equivalently written

$$\mathrm{KE}_n^{\mathrm{cla}}(\boldsymbol{\varphi}) = \frac{m_n}{2(n-1)} \sum_{j=1}^{n-1} \varphi_j^2 \tag{3.36}$$

with $\boldsymbol{\varphi}$ defined in Eq. (3.5), exhibits a *maximal* integrated autocorrelation time if the Cayley angle $\theta(x) = 2\arctan(x/2)$ is used. Indeed, the integrated autocorrelation time of this estimator is bounded by

$$\frac{\mathrm{aVar}(\mathrm{KE}_n^{\mathrm{cla}})}{\mathrm{Var}(\mathrm{KE}_n^{\mathrm{cla}})} \leq 1 + \frac{1}{2} \max_{1\leq j\leq n-1} \left| \frac{1 - A(\omega_{j,n}\Delta t)}{1 + A(\omega_{j,n}\Delta t)} \right| \, , \tag{3.37}$$

where the function $\big|\big(1 - A(x)\big)/\big(1 + A(x)\big)\big|$ monotonically increases as $\theta(x)$ approaches the largest stable angle for fixed x and α, which again corresponds to the Cayley angle.

We stress that the conclusions of this section hold for arbitrary friction schedules despite our use of the infinite-friction limit in Eqs. (3.35) and (3.37).

3.10 Appendix F: Stability interval calibration for liquid water simulations

This section describes the computational procedure used to identify $\Delta t = 1.4$ fs as close to the upper bound of the stability interval of T-RPMD applied to q-TIP4P/F liquid water at 298 K and 0.998 g/cm^3. Starting from an ensemble of 10^4 thermally distributed initial conditions obtained as per the following paragraph, we integrated an ensemble of 10^4 T-RPMD trajectories using the algorithm outlined in Sec. 3.2.1 in its single-bead realization (which is identical to velocity Verlet in classical MD [46]). We then counted the fraction of trajectories that did not "blow up," or exhibit numerical overflow, throughout their duration for each tested time-step. A time-step was deemed "stable" if 99% or more of the corresponding trajectories did not exhibit numerical overflow over a 50-ps time period. A range of time-steps between 1 ps and 2 ps was tested, and the fraction of stable trajectories at each time-step is reported in Fig. 3.8.

To avoid initialization bias in the stability interval estimation, thermalized initial phase-points were generated with a Metropolized Markov-chain Monte Carlo sampler targeted at the equilibrium configurational distribution of the liquid. Specifically, a randomized Hamiltonian Monte Carlo [9, 80] (rHMC) simulation of sufficient length was used to thermalize a crystalline configuration of the system at the target density, and 10^2 configurations were extracted from well-separated points along the rHMC trajectory. Each of these (approximately) independent draws from the equilibrium *configurational* distribution of the liquid at the target physical conditions was subsequently paired with 10^2 independent velocities drawn from the corresponding Maxwell–Boltzmann distribution, yielding 10^4 (approximately) independent draws from the *phase space* distribution of the liquid at thermal equilibrium.

3.11 Appendix G: Simulation and estimation details

This section compiles simulation protocols and statistical estimation methods used to generate Figs. 3.4 for the one-dimensional quantum harmonic oscillator, and Figs. 3.5 and 3.6 for room-temperature liquid water.

3.11.1 One-dimensional quantum harmonic oscillator

Numerical equilibrium averages and integrated autocorrelation times for the quantum harmonic oscillator were estimated by averaging over a 10-nanosecond T-RPMD trajectory integrated using the algorithm listed in Sec. 3.2.1, and initialized at an exact sample from the numerical stationary distribution (listed for the jth ring polymer mode in Eq. (3.19)) corresponding to the physical parameters (i.e., Λ, m, and β) and

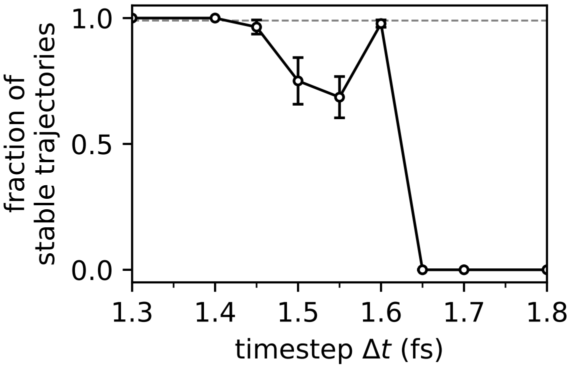

Figure 3.8. Stability interval calibration for q-TIP4P/F room-temperature liquid water simulations. Data points correspond to the fraction of thermally initialized single-bead T-RPMD trajectories that "survived" (i.e., did not exhibit numerical overflow) over a 50-ps time period at the respective integration time-step Δt. Error bars correspond to the standard error of the fraction of surviving trajectories across initialization points with different configurations. The gray dashed line marks the $\geq 99\%$ trajectory survival threshold beyond which we deem a time-step "stable." Note that no time-step beyond $\Delta t = 1.4$ fs reaches the survival threshold.

simulation parameters (i.e., n, Δt, and the function θ) listed in Sec. 3.3.1. Specifically, the statistics reported in Fig. 3.4 were obtained by partitioning the T-RPMD trajectory into 10 disjoint blocks, estimating the equilibrium average and autocorrelation time within each block, and computing the sample mean and standard error among the resulting block estimates with 1000 bootstrap resamples.

We now describe the formulas and methods used to obtain block estimates for the equilibrium mean and integrated autocorrelation time. The equilibrium average μ_{O_n} of observable O_n within each block of the partitioned T-RPMD trajectory was estimated using the standard estimator [81]

$$\hat{\mu}_{\mathsf{O}_n} = \frac{1}{K} \sum_{k=0}^{K-1} \mathsf{O}_n^{(k\Delta t)} \ , \tag{3.38}$$

where K is the number of steps in the block (i.e., the block size) and $\mathsf{O}_n^{(k\Delta t)}$ the value of O_n at the kth step within the block. Similarly, the lag-$k\Delta t$ autocovariance $C_{\mathsf{O}_n}(k\Delta t)$ was estimated using [81]

$$\hat{C}_{\mathsf{O}_n}(k\Delta t) = \sum_{\ell=0}^{K-k-1} \frac{\big(\mathsf{O}_n^{(\ell\Delta t)} - \hat{\mu}_{\mathsf{O}_n}\big)\big(\mathsf{O}_n^{((\ell+k)\Delta t)} - \hat{\mu}_{\mathsf{O}_n}\big)}{K-k}$$

for $0 \leq k\Delta t \leq (K-1)\Delta t = 1$ ns. The integrated autocorrelation time was subsequently estimated using [71, 81]

$$\frac{\widehat{\mathrm{aVar}_{\mathsf{O}_n}}}{\mathrm{Var}_{\mathsf{O}_n}}(M) = 1 + 2\sum_{k=1}^{M} \frac{\hat{C}_{\mathsf{O}_n}(k\Delta t)}{\hat{C}_{\mathsf{O}_n}(0)} \ , \tag{3.39}$$

where $0 < M \leq K$ is a suitable cutoff. The choice of M is nontrivial, as it carries a trade-off between bias (more pronounced at small M) and variance (more pronounced at large M) [71]. To choose M judiciously, we follow the *automatic windowing* (AW) method described in Appendix C of Ref. [82]. The AW method dictates that M should correspond to the smallest lag that satisfies the inequality

$$M \geq c\,\frac{\widehat{\mathrm{aVar}_{\mathsf{O}_n}}}{\mathrm{Var}_{\mathsf{O}_n}}(M) \ ,$$

where the parameter $c > 0$ dictates the variance-bias trade-off in place of M, and is chosen as large as possible to reduce the bias of the estimator for a given variance threshold.

Fig. 3.9 illustrates usage of the AW method for integrated autocorrelation time estimation, using trajectory data generated by the T-RPMD scheme with the choice $\theta(x) = 2\arctan(x/2)$ at $n = 64$ beads and $\Delta t = 2.0$ fs, and focusing on the observables $\mathrm{KE}_n^{\mathrm{pri}}$ (black), $\mathrm{KE}_n^{\mathrm{vir}}$ (red), and $\mathrm{KE}_n^{\mathrm{cla}}$ (cyan) introduced in Sec. 3.3.1. The estimated integrated autocorrelation times are plotted with solid lines in Fig. 3.9a for various values of c, and the corresponding cutoffs M are plotted in Fig. 3.9b. Exact integrated autocorrelation times are plotted with dashed lines in Fig. 3.9a. Note that as c (and thus M) increases, the estimates converge to the corresponding exact values at the expense of a larger variance, which can nonetheless be controlled by adjusting the block size K.

3.11.2 Room-temperature liquid water

The equilibrium averages and integrated autocorrelation times reported in Fig. 3.5 were obtained by averaging over 10-nanosecond T-RPMD trajectories integrated for each considered bead number n, time-step Δt, and function θ. All trajectories were initialized at an approximate sample from the corresponding numerical equilibrium distribution, obtained by thermalizing for 20 ps a classical (i.e., $n = 1$) configuration of the system into the n-bead ring-polymer phase space. Figure 3.10 below checks that this thermalization protocol indeed lead to near-equilibrium initialization of the T-RPMD trajectories. The reference equilibrium averages plotted with dashed lines

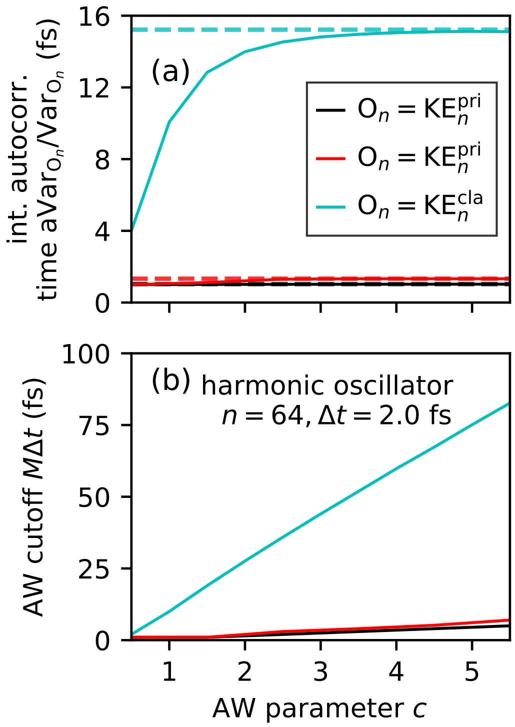

Figure 3.9. Integrated autocorrelation times of several observables of the one-dimensional harmonic oscillator in Sec. 3.3.1, estimated with the AW method. Trajectory data for the estimates was generated using the T-RPMD scheme with $\theta(x) = 2\arctan(x/2)$ at $n = 64$ beads and $\Delta t = 2.0$ fs, and processed as described in the current section. Estimated (resp. exact) integrated autocorrelation times for observables $\mathrm{KE}_n^{\mathrm{pri}}$ (black), $\mathrm{KE}_n^{\mathrm{vir}}$ (red), and $\mathrm{KE}_n^{\mathrm{cla}}$ (cyan) are shown in solid (resp. dashed) lines in panel a as a function of the windowing parameter c. Panel b plots the cutoffs determined by the choice of c for the three observables, where the linear relation between $M\Delta t$ and c at large values of the latter corroborates the non-spurious convergence of the autocorrelation time estimates.

in Fig. 3.5 were obtained by averaging over a one-nanosecond, 256-bead staging PIMD [11] trajectory integrated at a 0.1-fs time-step with the mass and friction parameters recommended in Ref. [8], and initialized with the same protocol used for the T-RPMD simulations.

The observables considered in Fig. 3.5 measure properties per H atom or per $\mathrm{H_2O}$ molecule, and thus the reported values are averages over estimates obtained for each simulated moiety. The equilibrium mean and integrated autocorrelation time of observable O_n for each moiety was estimated by partitioning the trajectory of the moiety into 10 disjoint 1-nanosecond blocks, evaluating Eqs. (3.38) and (3.39) within each block, and determining the sample mean and standard error among the block estimates with 1000 bootstrap resamples. The AW method [82] was applied to choose a cutoff lag $M \leq 1$ ns in Eq. (3.39), as illustrated in Fig. 3.9 for the harmonic oscillator application.

The T-RPMD trajectories used to generate Fig. 3.5 also yielded Fig. 3.6, where panels a and c plot autocovariance functions of the form $\frac{1}{N_{\mathrm{H_2O}}}\sum_{i=1}^{N_{\mathrm{H_2O}}}\left\langle \bar{\mathsf{O}}_i(0)\cdot\bar{\mathsf{O}}_i(k\Delta t)\right\rangle$, where $N_{\mathrm{H_2O}} = 32$ is the number of simulated $\mathrm{H_2O}$ molecules and $\bar{\mathsf{O}}_i(k\Delta t)$ is the bead-averaged value of observable O (e.g., the molecular dipole moment or center-of-mass velocity) on the ith molecule at time $k\Delta t$ along a stationary T-RPMD trajectory. The autocovariance $\left\langle \bar{\mathsf{O}}_i(0)\cdot\bar{\mathsf{O}}_i(t)\right\rangle$ was estimated for the lags $k\Delta t$ shown in Fig. 3.6 by

$$\left\langle \bar{\mathsf{O}}_i(0)\cdot\bar{\mathsf{O}}_i(k\Delta t)\right\rangle \approx \sum_{\ell=0}^{K-k-1} \frac{\bar{\mathsf{O}}_i^{(\ell\Delta t)}\cdot\bar{\mathsf{O}}_i^{((\ell+k)\Delta t)}}{K-k} ,$$

where $K\Delta t = 1$ ns is the length of each block in a partitioned 10-nanosecond T-RPMD trajectory. As with the results in Fig. 3.5, statistics for each molecule were obtained from block estimates via bootstrapping, and Figs. 3.6a and 3.6c report molecule-averaged statistics.

Figure 3.10 validates the 20-ps thermalization interval used to initialize the trajectories that generated Figs. 3.5 and 3.6. In detail, Figs. 3.10a and 3.10b (resp., Figs. 3.10c and 3.10d) plot the non-equilibrium mean of the primitive and virial quantum kinetic energy per H atom (resp. the mean $\mathrm{O-H}$ bond and $\mathrm{H-O-H}$ angle potential energy per water molecule) as it approaches the equilibrium value in Figs. 3.5a and 3.5c (resp., Figs. 3.5e and 3.5g) for a 64-bead ring polymer at a 1.4 fs time-step with the considered choices of θ. At each time $k\Delta t$ within the 20-ps time interval, the non-equilibrium mean is estimated by averaging across 1000 independent trajectories initialized at a point-mass distribution on the n-bead ring-polymer phase space centered at the classical (i.e., $n = 1$) sample used to initialize the reported simulations. Within statistical uncertainty, the non-equilibrium mean for each observable converges to its equilibrium value within the 20-ps interval at visually indistinguishable rates across the tested choices of θ.

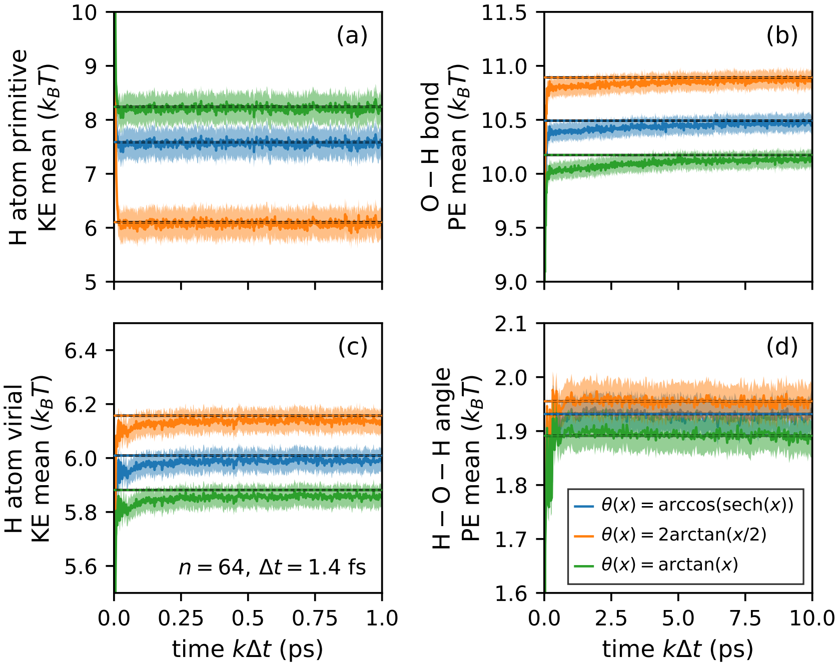

Figure 3.10. Convergence to equilibrium of the BAOAB-like schemes considered in Sec. 3.3.2 with $n = 64$ ring-polymer beads and a $\Delta t = 1.4$ fs time-step. With respect to the non-equilibrium 64-bead configurational distribution evolved from a point mass at a classical (i.e., $n = 1$) configuration, panels a and c plot the mean kinetic energy per H atom for the n-bead system as per the primitive and virial estimators, respectively, for times up to 1.0 ps. Panels b and d, respectively, plot the non-equilibrium mean O − H-bond and H − O − H-angle potential energy per q-TIP4P/F water molecule, for times up to 10 ps. The lightly shaded interval around each curve corresponds to the standard error of the estimated non-equilibrium mean, computed with 1000 bootstrap resamples from a sample of 1000 independent trajectories.

References Cited in Part I

[1] M. Parrinello and A. Rahman, “Study of an F center in molten KCl”, The Journal of Chemical Physics **80**, 860 (1984).

[2] G. J. Martyna, A. Hughes, and M. E. Tuckerman, “Molecular dynamics algorithms for path integrals at constant pressure”, The Journal of Chemical Physics **110**, 3275 (1999).

[3] P. Minary, G. J. Martyna, and M. E. Tuckerman, “Algorithms and novel applications based on the isokinetic ensemble. I. Biophysical and path integral molecular dynamics”, The Journal of Chemical Physics **118**, 2510 (2003).

[4] A. Beskos, G. Roberts, A. Stuart, and J. Voss, “MCMC methods for diffusion bridges”, Stochastics and Dynamics **8**, 319 (2008).

[5] A. Beskos, F. J. Pinski, J. M. Sanz-Serna, and A. M. Stuart, “Hybrid Monte Carlo on Hilbert spaces”, Stochastic Processes and their Applications **121**, 2201 (2011).

[6] J. Lu, Y. Lu, and Z. Zhou, “Continuum limit and preconditioned Langevin sampling of the path integral molecular dynamics”, Journal of Computational Physics **423**, 109788 (2020).

[7] Z. Zhang, X. Liu, Z. Chen, H. Zheng, K. Yan, and J. Liu, “A unified thermostat scheme for efficient configurational sampling for classical/quantum canonical ensembles via molecular dynamics”, The Journal of Chemical Physics **147**, 034109 (2017).

[8] J. Liu, D. Li, and X. Liu, “A simple and accurate algorithm for path integral molecular dynamics with the Langevin thermostat”, The Journal of Chemical Physics **145**, 024103 (2016).

[9] N. Bou-Rabee and J. M. Sanz-Serna, “Geometric integrators and the Hamiltonian Monte Carlo method”, Acta Numerica **27**, 113 (2018).

[10] N. Bou-Rabee and A. Eberle, “Two-scale coupling for preconditioned Hamiltonian Monte Carlo in infinite dimensions”, Stochastics and Partial Differential Equations: Analysis and Computations **9**, 207 (2021).

[11] M. E. Tuckerman, B. J. Berne, G. J. Martyna, and M. L. Klein, “Efficient molecular dynamics and hybrid Monte Carlo algorithms for path integrals”, The Journal of Chemical Physics **99**, 2796 (1993).

[12] I. R. Craig and D. E. Manolopoulos, “Quantum statistics and classical mechanics: Real time correlation functions from ring polymer molecular dynamics”, The Journal of Chemical Physics **121**, 3368 (2004).

[13] S. Habershon, D. E. Manolopoulos, T. E. Markland, and T. F. Miller III, “Ring-polymer molecular dynamics: Quantum effects in chemical dynamics from classical trajectories in an extended phase space”, Annual Reviews of Physical Chemistry **64**, 387 (2013).

[14] R. P. Feynman, *Quantum mechanics and path integrals* (McGraw-Hill, 1965).

[15] D. Chandler and P. G. Wolynes, “Exploiting the isomorphism between quantum theory and classical statistical mechanics of polyatomic fluids”, The Journal of Chemical Physics **74**, 4078 (1981).

[16] T. F. Miller III and D. E. Manolopoulos, “Quantum diffusion in liquid water from ring polymer molecular dynamics”, The Journal of Chemical Physics **123**, 154504 (2005).

[17] T. F. Miller III and D. E. Manolopoulos, “Quantum diffusion in liquid para-hydrogen from ring-polymer molecular dynamics”, The Journal of Chemical Physics **122**, 184503 (2005).

[18] M. Ceriotti, M. Parrinello, T. E. Markland, and D. E. Manolopoulos, “Efficient stochastic thermostatting of path integral molecular dynamics”, The Journal of Chemical Physics **133**, 124104 (2010).

[19] M. Ceriotti, D. E. Manolopoulos, and M. Parrinello, “Accelerating the convergence of path integral dynamics with a generalized Langevin equation”, The Journal of Chemical Physics **134**, 084104 (2011).

[20] M. Rossi, M. Ceriotti, and D. E. Manolopoulos, “How to remove the spurious resonances from ring polymer molecular dynamics”, The Journal of Chemical Physics **140**, 234116 (2014).

[21] M. Rossi, V. Kapil, and M. Ceriotti, “Fine tuning classical and quantum molecular dynamics using a generalized Langevin equation”, Journal of Chemical Physics **148**, 102301 (2018).

[22] G. Bussi, D. Donadio, and M. Parrinello, “Canonical sampling through velocity rescaling”, The Journal of Chemical Physics **126**, 014101 (2007).

[23] N. Bou-Rabee and H. Owhadi, “Long-run accuracy of variational integrators in the stochastic context”, SIAM Journal on Numerical Analysis **48**, 278 (2010).

[24] N. Bou-Rabee and E. Vanden-Eijnden, “Pathwise accuracy and ergodicity of metropolized integrators for SDEs”, Communications on Pure and Applied Mathematics **63**, 655 (2010).

[25] B. Leimkuhler and C. Matthews, “Rational construction of stochastic numerical methods for molecular sampling”, Applied Mathematics Research eXpress **2013**, 34 (2013).

[26] R. Korol, N. Bou-Rabee, and T. F. Miller III, “Cayley modification for strongly stable path-integral and ring-polymer molecular dynamics”, The Journal of Chemical Physics **151**, 124103 (2019).

[27] H. F. Trotter, “On the product of semi-groups of operators”, Proceedings of the American Mathematical Society **10**, 545 (1959).

[28] T. Matsubara, “A new approach to quantum-statistical mechanics”, Progress of Theoretical Physics **14**, 351 (1955).

[29] D. Li, X. Han, Y. Chai, C. Wang, Z. Zhang, Z. Chen, J. Liu, and J. Shao, “Stationary state distribution and efficiency analysis of the Langevin equation via real or virtual dynamics”, The Journal of Chemical Physics **147**, 184104 (2017).

[30] S. Kakutani, “On equivalence of infinite product measures”, Annals of Mathematics **49**, 214 (1948).

[31] V. I. Bogachev, *Gaussian measures*, Vol. 62, Mathematical Surveys and Monographs (American Mathematical Society, 1998).

[32] A. Pérez and M. E. Tuckerman, “Improving the convergence of closed and open path integral molecular dynamics via higher order Trotter factorization schemes”, The Journal of Chemical Physics **135**, 064104 (2011).

[33] O. Marsalek, P.-Y. Chen, R. Dupuis, M. Benoit, M. Méheut, Z. Bačić, and M. E. Tuckerman, “Efficient calculation of free energy differences associated with isotopic substitution using path-integral molecular dynamics”, Journal of Chemical Theory and Computation **10**, 1440 (2014).

[34] J. A. Barker, “A quantum-statistical Monte Carlo method; path integrals with boundary conditions”, The Journal of Chemical Physics **70**, 2914 (1979).

[35] M. F. Herman, E. J. Bruskin, and B. J. Berne, “On path integral Monte Carlo simulations”, The Journal of Chemical Physics **76**, 5150 (1982).

[36] B. García-Archilla, J. M. Sanz-Serna, and R. D. Skeel, “Long-time-step methods for oscillatory differential equations”, SIAM Journal on Scientific Computing **20**, 930 (1998).

[37] E. Hairer and C. Lubich, “Long-time energy conservation of numerical methods for oscillatory differential equations”, SIAM Journal on Numerical Analysis **38**, 414 (2000).

[38] J. M. Sanz-Serna, “Mollified impulse methods for highly oscillatory differential equations”, SIAM Journal on Scientific Computing **46**, 1040 (2008).

[39] R. I. McLachlan and A. Stern, “Modified trigonometric integrators”, SIAM Journal on Scientific Computing **52**, 1378 (2014).

[40] A. L. Gibbs and F. E. Su, “On choosing and bounding probability metrics”, International Statistical Review **70**, 419 (2002).

[41] S. Habershon, T. E. Markland, and D. E. Manolopoulos, “Competing quantum effects in the dynamics of a flexible water model”, The Journal of Chemical Physics **131**, 024501 (2009).

[42] S. Habershon, G. S. Fanourgakis, and D. E. Manolopoulos, “Comparison of path integral molecular dynamics methods for the infrared absorption spectrum of liquid water”, The Journal of Chemical Physics **129**, 074501 (2008).

[43] T. E. Markland and M. Ceriotti, “Nuclear quantum effects enter the mainstream”, Nature Reviews Chemistry **2**, 0109 (2018).

[44] D. Frenkel and B. Smit, *Understanding molecular simulation: From algorithms to applications*, 2nd (Academic Press, 2002).

[45] D. C. Rapaport, *The art of molecular dynamics simulation*, 2nd (Cambridge University Press, 2004).

[46] B. Leimkuhler and C. Matthews, *Molecular dynamics*, Vol. 39, Interdisciplinary Applied Mathematics (Springer International Publishing, 2015).

[47] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids*, 2nd edition (Oxford University Press, 2017).

[48] J. Cao and G. A. Voth, “The formulation of quantum statistical mechanics based on the Feynman path centroid density. II. Dynamical properties”, The Journal of Chemical Physics **100**, 5106 (1994).

[49] J. Liu, “Path integral Liouville dynamics for thermal equilibrium systems”, The Journal of Chemical Physics **140**, 224107 (2014).

[50] T. J. H. Hele, M. J. Willatt, A. Muolo, and S. C. Althorpe, “Boltzmann-conserving classical dynamics in quantum time-correlation functions: “Matsubara dynamics””, The Journal of Chemical Physics **142**, 134103 (2015).

[51] T. J. H. Hele, M. J. Willatt, A. Muolo, and S. C. Althorpe, “Communication: Relation of centroid molecular dynamics and ring-polymer molecular dynamics to exact quantum dynamics”, The Journal of Chemical Physics **142**, 191101 (2015).

[52] J. R. Cendagorta, Z. Bačić, and M. E. Tuckerman, “An open-chain imaginary-time path-integral sampling approach to the calculation of approximate symmetrized quantum time correlation functions”, The Journal of Chemical Physics **148**, 102340 (2018).

[53] I. R. Craig and D. E. Manolopoulos, “Chemical reaction rates from ring polymer molecular dynamics”, The Journal of Chemical Physics **122**, 084106 (2005).

[54] I. R. Craig and D. E. Manolopoulos, “A refined ring polymer molecular dynamics theory of chemical reaction rates”, The Journal of Chemical Physics **123**, 034102 (2005).

[55] A. Witt, S. D. Ivanov, M. Shiga, H. Forbert, and D. Marx, “On the applicability of centroid and ring polymer path integral molecular dynamics for vibrational spectroscopy”, The Journal of Chemical Physics **130**, 194510 (2009).

[56] J. A. Morrone and R. Car, “Nuclear quantum effects in water”, Physical Review Letters **101**, 017801 (2008).

[57] B. Cheng, E. A. Engel, J. Behler, C. Dellago, and M. Ceriotti, “Ab initio thermodynamics of liquid and solid water”, Proceedings of the National Academy of Sciences **116**, 1110 (2019).

[58] T. Zimmermann and J. Vaníček, “Path integral evaluation of equilibrium isotope effects”, The Journal of Chemical Physics **131**, 024111 (2009).

[59] D. L. Eldridge, R. Korol, M. K. Lloyd, A. C. Turner, M. A. Webb, T. F. Miller III, and D. A. Stolper, “Comparison of experimental vs theoretical abundances of $^{13}CH_3D$ and $^{12}CH_2D_2$ for isotopically equilibrated systems from 1 to 500 °C”, ACS Earth and Space Chemistry **3**, 2747 (2019).

[60] G. Strang, “On the construction and comparison of difference schemes”, SIAM Journal on Numerical Analysis **5**, 506 (1968).

[61] M. P. Calvo and J. M. Sanz-Serna, “Instabilities and inaccuracies in the integration of highly oscillatory problems”, SIAM Journal on Scientific Computing **31**, 1653 (2009).

[62] V. I. Arnol’d, *Mathematical methods of classical mechanics*, 2nd edition, Vol. 60, Graduate Texts in Mathematics (Springer Science & Business Media, 2013).

[63] B. J. Braams and D. E. Manolopoulos, “On the short-time limit of ring polymer molecular dynamics”, The Journal of Chemical Physics **125**, 124105 (2006).

[64] N. Bou-Rabee, “Cayley splitting for second-order Langevin stochastic partial differential equations”, Preprint (2017).

[65] G. Metafune, D. Pallara, and E. Priola, “Spectrum of Ornstein-Uhlenbeck operators in L^p spaces with respect to invariant measures”, Journal of Functional Analysis **196**, 40 (2002).

[66] G. A. Pavliotis, *Stochastic Processes and Applications*, Vol. 60, Texts in Applied Mathematics (Springer, 2014).

[67] J. M. Sanz-Serna and M. P. Calvo, *Numerical Hamiltonian problems*, Vol. 7, Applied mathematics and mathematical computation (Chapman & Hall, 1994).

[68] C. Villani, *Optimal transport: Old and new*, Vol. 338, Comprehensive Studies in Mathematics (Springer Science & Business Media, 2008).

[69] C. R. Givens and R. M. Shortt, “A class of Wasserstein metrics for probability distributions”, Michigan Mathematical Journal **31**, 231 (1984).

[70] C. J. Geyer, “Practical Markov chain Monte Carlo”, Statistical Science **7**, 473 (1992).

[71] A. Sokal, “Monte Carlo methods in statistical mechanics: Foundations and new algorithms”, in *Functional integration: Basics and applications*, edited by C. DeWitt-Morette, P. Cartier, and A. Folacci (Springer Science & Business Media, 1997), pp. 131–192.

[72] S. Asmussen and P. W. Glynn, *Stochastic simulation: Algorithms and analysis*, Vol. 57, Stochastic Modelling and Applied Probability (Springer, 2007).

[73] R. D. Skeel and Y. Fang, “Comparing Markov chain samplers for molecular simulation”, Entropy **19**, 561 (2017).

[74] Y. Fang, Y. Cao, and R. D. Skeel, “Quasi-reliable estimates of effective sample size”, preprint (2017).

[75] R. Korol, J. L. Rosa-Raíces, N. Bou-Rabee, and T. F. Miller III, “Dimension-free path-integral molecular dynamics without preconditioning”, The Journal of Chemical Physics **152**, 104102 (2020).

[76] Y. V. Suleimanov, J. W. Allen, and W. H. Green, “RPMDrate: Bimolecular chemical reaction rates from ring polymer molecular dynamics”, Computer Physics Communications **184**, 833 (2013).

[77] V. Kapil, M. Rossi, O. Marsalek, R. Petraglia, Y. Litman, T. Spura, B. Cheng, A. Cuzzocrea, R. H. Meißner, D. M. Wilkins, B. A. Helfrecht, P. Juda, S. P. Bienvenue, W. Fang, J. Kessler, I. Poltavsky, S. Vandenbrande, J. Wieme, C. Corminboeuf, T. D. Kühne, D. E. Manolopoulos, T. E. Markland, J. O. Richardson, A. Tkatchenko, G. A. Tribello, V. Van Speybroeck, and M. Ceriotti, “i-PI 2.0: A universal force engine for advanced molecular simulations”, Computer Physics Communications **236**, 214 (2019).

[78] M. Abramowitz and I. A. Stegun, *Handbook of mathematical functions, with formulas, graphs, and mathematical tables* (Dover Publications, Inc., 1965).

[79] T. Andreescu, *Essential linear algebra with applications* (Birkhäuser Basel, 2016).

[80] N. Bou-Rabee and J. M. Sanz-Serna, “Randomized Hamiltonian Monte Carlo”, Annals of Applied Probability **27**, 2159 (2017).

[81] M. B. Priestly, *Spectral analysis and time series* (Academic Press, 1981).

[82] N. Madras and A. Sokal, “The pivot algorithm: A highly efficient Monte Carlo method for the self-avoiding walk”, Journal of Statistical Physics **50**, 109 (1988).

Part II

PARALLEL-IN-TIME STOCHASTIC MOLECULAR DYNAMICS

Chapter 4

Path-accelerated stochastic molecular dynamics: Parallel-in-time integration using path integrals

Adapted from

J. L. Rosa-Raíces, B. Zhang, and T. F. Miller III, "Path-accelerated stochastic molecular dynamics: Parallel-in-time integration using path integrals", The Journal of Chemical Physics **151**, 164120 (2019) DOI: `10.1063/1.5125455`.

Abstract

Massively parallel computer architectures create new opportunities for the performance of long-timescale stochastic molecular dynamics (MD) simulations. Here, we introduce the path-accelerated molecular dynamics (PAMD) method that takes advantage of distributed computing to reduce the wall-clock time of MD simulation via parallelization with respect to stochastic MD time-steps. The marginal distribution for the time evolution of a system is expressed in terms of a path integral, enabling the use of path sampling techniques to numerically integrate MD trajectories. By parallelizing the evaluation of the path action with respect to time and by initializing the path configurations from a non-equilibrium distribution, the algorithm enables significant speedups in terms of the length of MD trajectories that can be integrated in a given amount of wall-clock time. The method is demonstrated for Brownian dynamics, although it is generalizable to other stochastic equations of motion. We apply the method to two simple systems, a harmonic oscillator and a Lennard–Jones liquid, and we show that in comparison to the conventional Euler integration scheme for Brownian dynamics, the new method can reduce the wall-clock time for integrating trajectories of a given length by more than three orders of magnitude in the former system and more than two in the latter. This new method for parallelizing MD in the dimension of time can be trivially combined with algorithms for parallelizing the MD force evaluation to achieve further speedup.

4.1 Introduction

Molecular dynamics (MD) [1, 2] is the central tool for simulating chemical, biological, and materials systems, with new algorithms and hardware expanding the range of accessible scales in space and in time [3–5]. Faster processors have played an important role in this expansion, although the most dramatic improvements in recent years have come from the number of available processors, rather than the clock-speed of the individual cores [6, 7]. In particular, highly multi-threaded computer architectures have been used to parallelize the MD force evaluation, greatly reducing the wall-clock time needed to perform an individual MD step [8–13]. However, despite this progress in the parallelization of MD simulations with respect to the force evaluations (i.e., in *space*), less attention has been dedicated to the notion of parallelization with respect to the MD time-steps (i.e., in *time*).

The sequential nature of MD (i.e., the need to have access to a given time-step before the next time-step can be computed) would seem to discount the possibility of exploiting parallelization in time; nonetheless, methods for parallel-in-time integration are being developed and applied to MD simulation. Most approaches [14–17] are based on a prediction-correction paradigm that combines fine (i.e., accurate and expensive) and coarse (i.e., inaccurate and inexpensive) solvers to iteratively refine approximations of a trajectory in a convergent and parallel-in-time fashion. A range of coarse solvers and iteration schemes have been employed to evaluate MD trajectories of molecular systems with parallelization in the time domain [18–23], leading to order-of-magnitude reductions in the wall-clock time-to-solution with respect to sequential integration at the fine level of accuracy. Schemes for approximate long-timescale integration via trajectory splicing are an alternative route to parallelization in time, yielding accurate time evolution for systems that exhibit strong timescale separation on well-characterized regions of the potential energy landscape [24].

The current chapter takes a different approach to parallelizing MD in time, focusing on *stochastic MD* equations of motion for systems in contact with a thermal bath. We demonstrate that by working with ensembles of trajectories in a path-integral framework, multiple processors can be employed to reduce the wall-clock time needed to evolve a stochastic MD trajectory of arbitrary length, without resorting to parallelization of the MD force evaluation. This method of parallelization for MD trajectories is independent of, and thus entirely complementary to, parallelization of the MD force evaluations, and it creates new opportunities to harness large numbers of available computer processors for the generation of long-timescale MD trajectories.

4.2 Method

4.2.1 MD integration based on path distributions

While our approach generalizes to other stochastic equations of motion, we illustrate it in the current chapter for Brownian (i.e., overdamped Langevin) dynamics under potential V at temperature β^{-1},

$$\dot{x}(t) = -\gamma^{-1} V'(x(t)) + \sqrt{2D}\, \dot{w}(t) \; , \tag{4.1}$$

where the diffusion coefficient D and the friction coefficient γ are related by the Einstein relation $D = (\beta\gamma)^{-1}$, and where $w(t)$ is a standard Wiener process. MD trajectories can be generated by discretizing Eq. (4.1) with various numerical integration schemes [25–29], such as the forward Euler algorithm [2]

$$x(t+dt) - x(t) = -\gamma^{-1} V'(x(t))\, dt + \sqrt{2Ddt}\, \xi \; , \tag{4.2}$$

where dt is the discretization time-step and ξ is a standard Gaussian random variate. The marginal distribution associated with time evolution of the system by dt according to Eq. (4.2) is [30]

$$K(x(t+dt)|x(t);dt) \propto \\ \exp\left\{ -\frac{dt}{4D} \left(\frac{x(t+dt)-x(t)}{dt} + \frac{V'(x(t))}{\gamma} \right)^2 \right\} \; , \tag{4.3}$$

such that the likelihood of a MD trajectory of length $T = N\,dt$ that evolves the system along positions $\boldsymbol{X} = \{x(t_0), x(t_1), \ldots, x(t_N)\}$ at times $t_n = t + n\,dt$ is

$$\prod_{n=0}^{N-1} K(x(t_{n+1})|x(t_n);dt) \equiv e^{-S_N[\boldsymbol{X}]} \; , \tag{4.4}$$

where $S_N[\boldsymbol{X}]$ is the action associated with the MD trajectory and, for the case of Eq. (4.1), is equivalent to the Onsager–Machlup action functional in a suitably defined continuous-time limit [31, 32].

From Eq. (4.4), the position of the time-evolved system at time T has a marginal distribution given by the path integral

$$K(x(T+t)|x(t);T) \propto \int_{\mathbb{R}} \mathrm{d}x_1 \cdots \int_{\mathbb{R}} \mathrm{d}x_{N-1}\, e^{-S_N[\boldsymbol{X}]} \; , \tag{4.5}$$

where $x_n = x(t_n)$. It is clear that this path-integral formulation of the ensemble of MD trajectories provides an equivalent description of the time evolution of the system as Eq. (4.2); for this reason, trajectories generated using the scheme introduced in this section are correctly referred to as MD trajectories. Numerous studies have explored path-based formulations of MD, with variations of the underlying equation of motion and of the discretization of the action [33–45].

Setting aside issues of efficiency until Sec. 4.2.2, we note that the path-integral formulation of the marginal distribution for the time-evolved system offers a simple MD integration scheme, illustrated in Fig. 4.1. First, *sampling* from the distribution of paths of length T, with likelihood given by Eq. (4.4), is performed using Monte Carlo (MC) or related methods (Fig. 4.1A) [46–48]; by drawing a random realization from this distribution, we obtain a segment of MD trajectory from time 0 to time T (illustrated by the heavy orange path in Fig. 4.1A). Then, by *shifting* from $x(0)$ to $x(T)$ along the sampled path, we resolve a trajectory from $x(0)$ to $x(T)$ (represented by the heavy green path in Fig. 4.1B) that is statistically equivalent to a realization from the Euler algorithm defined in Eq. (4.2). After shifting the left endpoint of the path from $x(0)$ to $x(T)$, we restart the path sampling to extend the trajectory from time T to time $2T$. Iteration of this scheme will lead to the numerical integration of a MD trajectory of arbitrary length in time.

Figure 4.2 illustrates a generalization of the integration scheme presented in Fig. 4.1. Figure 4.2A repeats Fig. 4.1A; we first sample a path of length T that is discretized into N_{path} time-steps (where $N_{\text{path}} = T/dt$) to obtain a realization of the path that is consistent with the marginal distribution of the time-evolved system for each time $\Delta t \leq T$. Then, in Fig. 4.2B, we shift the left endpoint of the sampled path (indicated in orange) by N_{shift} time-steps (where $N_{\text{shift}} = \Delta t/dt$) to the position $x(\Delta t)$. With the remaining segment of the path now located at positions $\{x(\Delta t), x(\Delta t + dt), \ldots, x(T)\}$, we grow the path out of $x(T)$ by N_{shift} time-steps to *regenerate* the original number of time-steps in the path. The positions of the system at the regenerated time-steps can be drawn from any distribution (and in Fig. 4.2B they are obtained via straight-line extrapolation). Finally, as illustrated in Fig. 4.2C, sampling is again performed to generate a path consistent with evolution from time Δt to time $T + \Delta t$; this sampling removes any artifacts introduced by the arbitrary distribution used to grow the shifted path. As for the scheme in Fig. 4.1, iterating through the sampling, shifting, and regeneration steps of the scheme in Fig. 4.2 yields a MD trajectory of arbitrary length in time that is statistically equivalent to a realization from

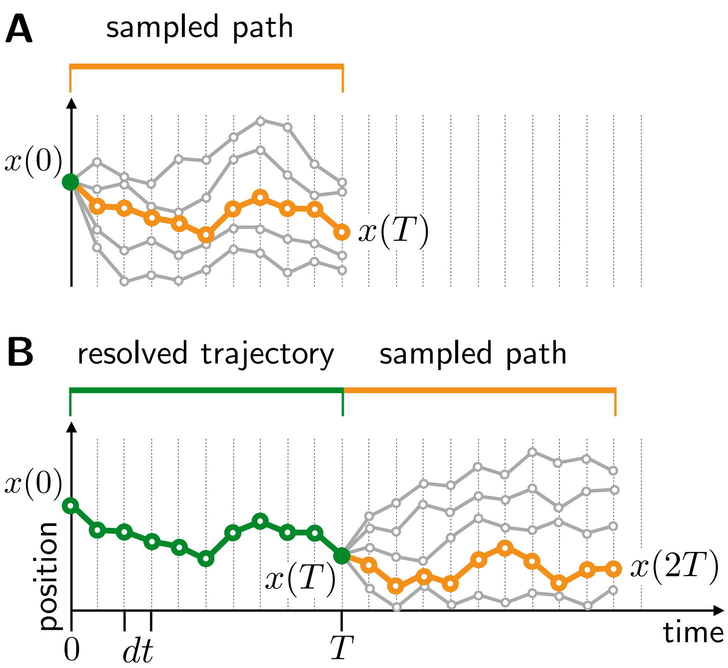

Figure 4.1. Illustration of a simple path-based MD integration scheme. **A** Sampling of the distribution of paths for the period of time from 0 to T. A particular path drawn from this distribution is indicated in orange. **B** Shifting along a sampled path (orange path in **A**) from $x(0)$ to $x(T)$, thereby resolving the segment of MD trajectory indicated in green, and then resuming the sampling of the paths for the period of time from T to $2T$.

the Euler algorithm. The only difference between these two path-based integration schemes is that Fig. 4.1 involves shifting along the full length of the sampled path, whereas Fig. 4.2 involves shifting only a fraction of the way along the sampled path.

Just like the Euler scheme in Eq. (4.2), the schemes illustrated in Figs. 4.1 and 4.2 enable the numerical integration of MD trajectories. Each of these integration schemes consist of sequential iterations of an elementary step that predicts the state of the system at some later time. In the Euler scheme, the prediction can be conducted analytically based on the distribution defined in Eq. (4.3). In path-based integration schemes employing path lengths longer that dt, however, no such analytical expression exists for general systems; path sampling is therefore needed before each shifting event to generate time-evolved system positions consistent with the correct marginal distribution.

For the scheme in Fig. 4.2, it is assumed that the path distributions in parts A and C are well sampled. For MC path sampling algorithms, this implies that the number of configurations of the path that are sampled in parts A and C, N_{sample}, is large

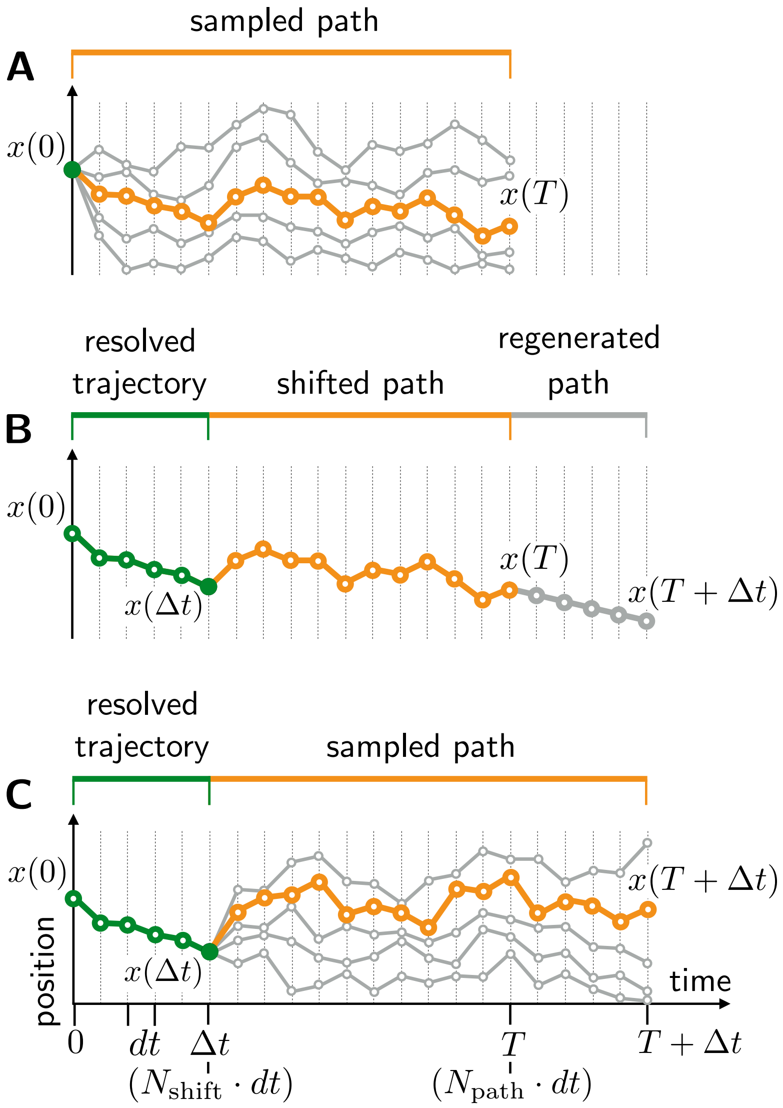

Figure 4.2. Illustration of a path-accelerated molecular dynamics (PAMD). **A** Sampling of the distribution of paths for the period of time from 0 to T. A particular path drawn from this distribution is indicated in orange. **B** Shifting along a sampled path (orange path in **A**) from $x(0)$ to $x(\Delta t)$, thereby resolving the segment of MD trajectory indicated in green, and regenerating the full length of the path by drawing positions for the system from time $T+dt$ to $T+\Delta t$ from an arbitrary distribution. **C** Sampling of the distribution of paths for the period of time from Δt to $\Delta t+T$.

in comparison to the number that is needed to generate uncorrelated realizations of the path. If N_{sample} is smaller than this decorrelation number, then the distribution of paths that is generated in part C may be biased by the way in which the path was regenerated in part B. However, the only requirement for generating accurate MD trajectories using the scheme in Fig. 4.2 is accurate sampling of paths consistent with the marginal distribution $K(x(\Delta t)|x(0);\Delta t)$; it is not essential that the marginal distribution associated with the full path, $K(x(T)|x(0);T)$, be sampled without error. Recalling that $T = N_{\mathrm{path}}\,dt$ and $\Delta t = N_{\mathrm{shift}}\,dt$, this suggests that for a given path sampling algorithm, there is an interplay between parameters N_{sample}, N_{path} and N_{shift}; for given values of N_{path} and N_{shift}, there is an associated number of path configurations (N_{sample}) that must be sampled in order to generate a sufficiently accurate marginal distribution $K(x(\Delta t)|x(0);\Delta t)$.

This interplay between N_{sample}, N_{path}, and N_{shift} is illustrated in Figs. 4.3A-C, which plot the error in the marginal distribution generated using the scheme in Fig. 4.2 for the Brownian dynamics of a harmonic oscillator. Full calculation details are provided in Sec. 4.3. The error plotted in Figs. 4.3A-C corresponds to the Kullback–Leibler divergence,

$$D_{\mathrm{KL}}(t) = \left\langle \int_{\mathbb{R}} \mathrm{d}x_t\, P(x_t|x_0;t) \log \frac{P(x_t|x_0;t)}{Q(x_t|x_0;t)} \right\rangle , \tag{4.6}$$

where $Q(x_t|x_0;t)$ is the marginal distribution estimated using sampled paths from the scheme in Fig. 4.2, and $P(x_t|x_0;t)$ is the exact marginal distribution. The angled brackets denote averaging with respect to the Boltzmann distribution of positions that is sampled by the exact dynamics, $P(x_0) = Z^{-1}e^{-\beta V(x_0)}$, where $Z = \int_{\mathbb{R}} \mathrm{d}x_0\, e^{-\beta V(x_0)}$ is the partition function. For a harmonic oscillator with potential $V(x) = \frac{1}{2}kx^2$ [49],

$$P(x_t|x_0;t) \propto \exp\left\{ -\frac{\beta k}{2} \frac{(x_t - e^{-\gamma^{-1}kt}x_0)^2}{(1 - e^{-2\gamma^{-1}kt})} \right\} , \tag{4.7}$$

and we employ $k = 1$ for the oscillator force constant, $\beta^{-1} = 1$ for the temperature and $\gamma = 1$ for the friction coefficient. $D_{\mathrm{KL}}(t)$ returns non-negative values that approach 0 as $Q(x_t|x_0;t)$ more accurately reproduces $P(x_t|x_0;t)$. As a function of time t along the sampled paths, $D_{\mathrm{KL}}(t)$ is plotted in Figs. 4.3A-C for seven simulations that employ the scheme in Fig. 4.2 with different values of N_{sample}, N_{path}, and N_{shift}. The results correspond to sampled paths of length $T \geq 1$ that are discretized into time-steps of $dt = 1/32$, and $D_{\mathrm{KL}}(t)$ is evaluated for the numerically generated marginal distributions at times $dt \leq t \leq 1$.

Figure 4.3A addresses the case where N_{sample} and N_{path} are held fixed and various values of N_{shift} are used. Comparison of the blue ($N_{\mathrm{shift}} = 4$), green ($N_{\mathrm{shift}} = 8$), and red ($N_{\mathrm{shift}} = 16$) curves shows that for a given value of N_{sample}, smaller values of N_{shift} lead to smaller errors in the numerically generated marginal distribution. Using the scheme in Fig. 4.2, a given segment of the path is sampled

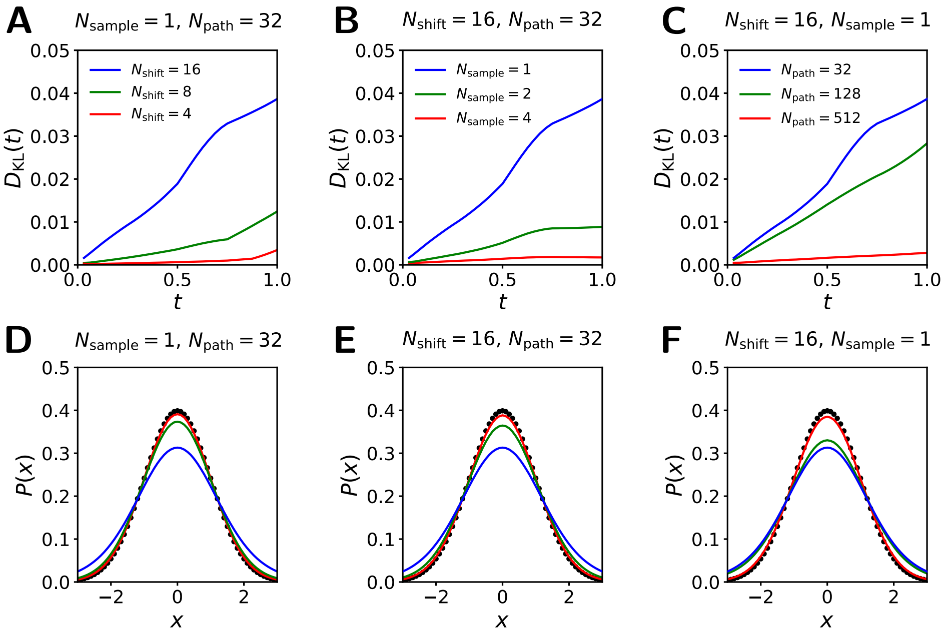

Figure 4.3. The interplay of the parameters N_{sample}, N_{path}, and N_{shift} in determining the accuracy of the PAMD integration scheme. Panels **A**, **B** and **C** show the Kullback–Leibler divergence $D_{\text{KL}}(t)$ defined in Eq. (4.6), from the marginal distribution generated using the scheme in Fig. 4.2 to the exact marginal distribution for the dynamics of an overdamped harmonic oscillator. Panels **D**, **E**, and **F** plot the position distributions $P(x)$ sampled by the MD trajectories integrated using the numerically generated marginal distributions from panels **A**, **B**, and **C**, respectively, in comparison to the exact position distribution shown in black dots. It is seen that the accuracy of the numerically generated marginal distribution dictates that of the integrated MD trajectory, and improves with decreasing N_{shift} (panels **A** and **D**), increasing N_{sample} (panels **B** and **E**), or increasing N_{path} (panels **C** and **F**) for given values of the remaining parameters (indicated at the top of each panel).

$N_{\text{path}} \cdot N_{\text{sample}}/N_{\text{shift}}$ times before it is used to generate the marginal distribution for the integration of the MD trajectory; therefore, smaller values of N_{shift} lead to better sampling of the path distribution and smaller errors in the marginal distribution.

Figure 4.3B illustrates a second scenario where N_{shift} and N_{path} are held fixed and increasing values of N_{sample} are used. Comparison of the blue ($N_{\text{sample}} = 1$), green ($N_{\text{sample}} = 2$), and red ($N_{\text{sample}} = 4$) curves shows that for a given value of N_{shift}, larger values of N_{sample} (i.e., more sampling per shifting event) lead to smaller errors in the numerically generated marginal distribution for integrating the MD trajectory. This result is intuitive, as more sampling leads to elimination of the bias associated with the arbitrary distribution used in the regeneration of the full length of the path.

In Fig. 4.3C, N_{shift} and N_{sample} are held fixed as the length of the sampled path (N_{path}) is increased while keeping the path discretization time-step unchanged. Comparison of the blue ($N_{\text{path}} = 32$), green ($N_{\text{path}} = 128$), and red ($N_{\text{path}} = 512$) curves demonstrates that increasing the total length $N_{\text{path}} \cdot dt$ of the sampled paths improves the accuracy of the numerically generated marginal distribution. Like decreasing N_{shift} for a given N_{path} (as in Fig. 4.3A), increasing N_{path} for a given N_{shift} allows for more sampling of each segment of the path employed to generate the marginal distribution associated with the MD time evolution.

While Figs. 4.3A-C illustrate the errors in the marginal distribution generated using the scheme in Fig. 4.2, Figs. 4.3D-F illustrate the corresponding errors in the equilibrium distribution that is sampled by the integrated MD trajectories. For the various employed parameters, the results from path-based MD integration are compared to the exact Boltzmann distribution (dots) and, as expected, the errors in the marginal distribution with given values of N_{sample}, N_{path}, and N_{shift} are reflected in the distribution of positions that are visited in the MD trajectories. Supplementary movies M1 and M2 (see Supplementary Material at the end of this chapter), respectively, illustrate the integration scheme in Fig. 4.2 as it generates the position distributions for the cases $N_{\text{sample}} = 1$ (blue curve) and $N_{\text{sample}} = 4$ (red curve) in Fig. 4.3E.

In summary, Fig. 4.3 demonstrates that decreasing N_{shift}, increasing N_{sample}, or increasing N_{path} leads to greater accuracy in the integrated MD trajectories; as will be shown in Sec. 4.2.2, the interplay between these three parameters is also critical for determining the computational efficiency of MD integration using the scheme in Fig. 4.2.

Before addressing efficiency, however, Fig. 4.4 illustrates that the integration scheme in Fig. 4.2 is a non-equilibrium relaxation process for the segments of the sampled path. For the case of the harmonic oscillator, Figs. 4.3A-C indicate that errors in the numerically generated marginal distributions are typically larger at the right endpoint of the sampled path. This trend emerges because the integration scheme regenerates path segments in configurations that are out of equilibrium (Fig. 4.2B). The light gray paths in Fig. 4.4 correspond to independent realizations of the sampled path (orange) obtained while integrating a harmonic oscillator trajectory. Since the segment at the nose of the path has undergone little sampling after regeneration, it is far from equilibrium with respect to the distribution of segments of an equilibrium harmonic oscillator trajectory (Fig. 4.4A; distributions at right). However, as

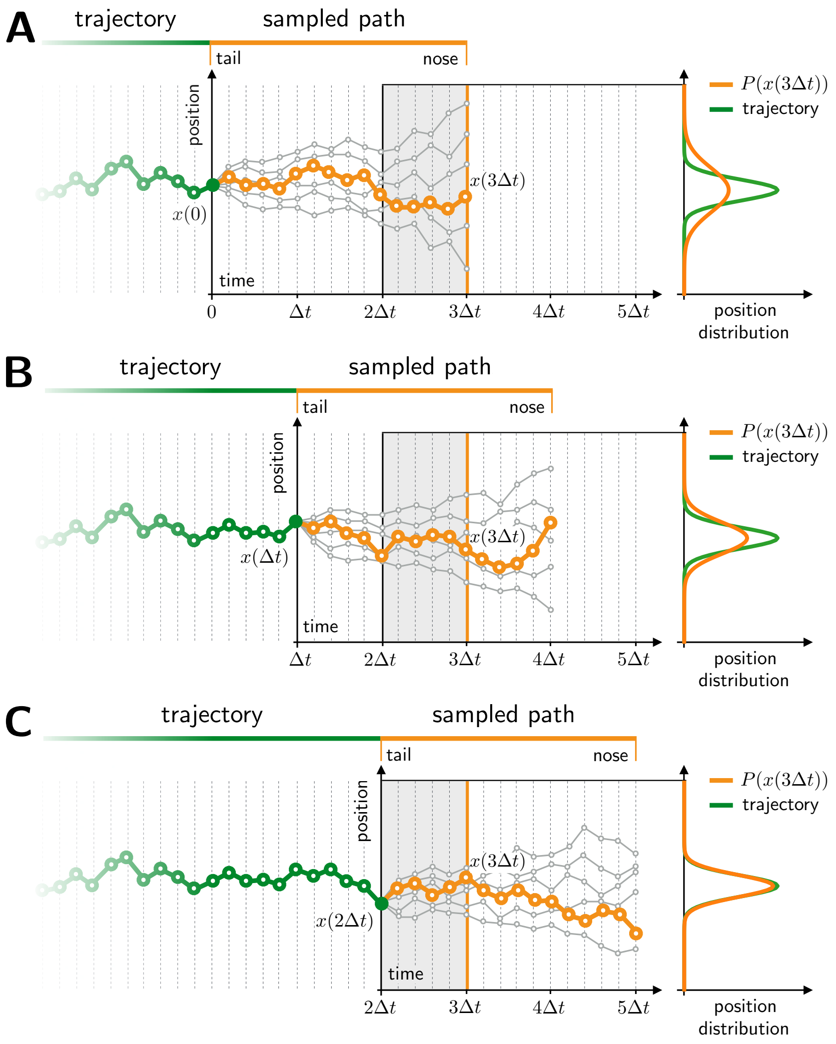

Figure 4.4. PAMD integrates equilibrium trajectories by relaxing non-equilibrium path segments. The state of the sampled path is shown for three consecutive iterations of the path-based integration scheme applied to a Brownian harmonic oscillator on the left of panels **A**, **B**, and **C**. The gray box in each panel highlights the configuration of a particular segment of the path after each iteration of the integration scheme. On the right of each panel are plotted the distributions $P(x(3\Delta t))$ of positions $x(3\Delta t)$ sampled by the boxed path segment at time $3\Delta t$ (orange curves), along with the distribution of configurations sampled by the integrated trajectory (green curve). The path segment inside the gray box in each panel reaches equilibrium by undergoing sampling as it shifts from the nose (right endpoint) to the tail (left endpoint) of the sampled path; accordingly, the distribution of sampled positions approaches that sampled by the integrated trajectory.

that segment works its way from the nose (right endpoint) to the tail (left endpoint) of the path, it is sampled with increasing accuracy (Figs. 4.4B and C). This relaxation process is illustrated by the distribution of positions, $P(x(3\Delta t))$, sampled by the foremost end $x(3\Delta t)$ of the path segment in the gray box throughout Fig. 4.4; as the orange curves indicate, this distribution approaches that sampled by the harmonic oscillator trajectory (green curve) as the segment relaxes toward equilibrium and simultaneously approaches the tail of the sampled path.

4.2.2 An opportunity for speedup

At face value, the path-based integration scheme in Fig. 4.2 may appear to be inefficient, given the difficulties of sampling uncorrelated paths [38, 46, 50, 51]. Yet, it has several potential advantages: Firstly, there is an opportunity for parallelization, given that typical expressions for the path action incur a dominant source of computational cost from the evaluation of the forces in the system along the path ($V'(x)$ in Eq. (4.3)). These forces can be evaluated independently, enabling straightforward parallelization of the action with respect to time. Secondly, regeneration of the path to its full length following shifting (Fig. 4.2B) can be performed using an arbitrary distribution to obtain the system positions for the regenerated time-steps; consequently, it is possible to carry out this operation at a cost that is negligible relative to evaluation of the MD forces. Thirdly, MC path sampling provides a numerically more stable way for generating trajectories than integration of the discretized equations of motion [34]; thus, a sufficiently accurate MD trajectory may be obtained with the path-based integration scheme at a larger time-step than a conventional Brownian dynamics integrator would allow.

The above considerations suggest that the scheme in Fig. 4.2 could lead to reduction of the wall-clock time associated with MD integration, in comparison with standard methods. To quantify the speedup achieved with the new scheme, we introduce a measure χ, with χ^{-1} defined as the number of force evaluations per processor per step of time dt_{E}, where dt_{E} is the time-step used by the Euler algorithm to integrate the Brownian dynamics. The wall-clock speedup of the path-based integration scheme is thus χ, assuming that (*i*) evaluation of the MD forces dominates the cost of the evaluation of the path action, (*ii*) parallel computer processors are used to independently evaluate the forces along the discretized path, and (*iii*) regeneration of the full length of the path following shifting (Fig. 4.2B) is performed without evaluating the MD forces. It is clear that for the Euler algorithm, $\chi = 1$, such that this measure provides a simple basis of comparison of the wall-clock time for the proposed path-based integration scheme (which employs parallelization in time) versus the wall-clock time for a conventional stochastic MD integration scheme (which does not). In the current chapter, we set aside the complementary issue of speeding up MD integration via parallelization within the force evaluation.

For a general implementation of the integration scheme in Fig. 4.2, the expression for χ is obtained as follows. Recalling previously introduced notation, we employ sampled paths of length T that are discretized with a time-step of dt, which may be different (and is typically larger [34]) than the numerically stable time-step for the Euler algorithm, dt_{E}. Let N_{force} be the number of MD force evaluations that are required during path sampling per shifting event, which depends on both N_{sample} and the details of the path sampling algorithm, and let N_{procs} be the number of employed parallel processors. Since the number of force evaluations per processor per shifting event is given by $N_{\mathrm{force}}/N_{\mathrm{procs}}$, the speedup is

$$\chi = N_{\mathrm{shift}} \cdot \frac{N_{\mathrm{procs}}}{N_{\mathrm{force}}} \cdot \frac{dt}{dt_{\mathrm{E}}} \; . \tag{4.8}$$

Equation (4.8) shows that the path-based integration scheme in Fig. 4.2 offers the possibility for reduction of the wall-clock time needed to compute MD trajectories, relative to conventional stochastic MD. Factors that enable this speedup include the increase in the path discretization time-step (dt) relative to that possible for conventional stochastic MD (dt_{E}), maximization of the number of integrated time-steps per shifting event (N_{shift}), maximization of the number of parallel processors to perform the independent force evaluations associated with the calculation of the path action (N_{procs}), and minimization of the number of force evaluations needed per shifting event (N_{force}). As will be shown in Sec. 4.4, this approach indeed enables substantial speedups in the integration of MD trajectories while preserving the accuracy of the dynamics, and we henceforth refer to the method as path-accelerated molecular dynamics (PAMD).

4.3 Calculation details

We implement the PAMD method with sampling of the path distribution through the multilevel *sliding and sampling* (S&S) algorithm introduced in Ref. [38]. For a path of N_{path} time-steps, a total of $L = \log_2 N_{\mathrm{path}}$ levels are defined (Fig. 4.5A); finer levels (smaller values of the level index $1 \leq l \leq L$) correspond to partitions of the path into fragments of increasingly smaller length where the local configuration of the path is sampled. In accordance with the S&S algorithm, neighboring path fragments share endpoints that are chosen randomly such that the length of the fragments varies from 1 to 2^l time-steps; we call this random fragmentation. For all fragmentations of the path at level l, internal fragments of the path are of length 2^l and fragments at the termini of the path have a combined length of 2^l. During a MC step for a given fragmentation of the path (Fig. 4.5B), the system positions at shared endpoints of neighboring path fragments are held fixed to permit mutually independent updates of the fragment configurations. Furthermore, the position of the system at the tail endpoint of the path is always fixed throughout the MC step, whereas that at the nose endpoint of the path undergoes sampling together with the nose fragment. Path fragment configurations are updated through application of the standard Metropolis–Hastings criterion [52, 53], with trial configurations drawn from a distribution that satisfies the boundary conditions at the fragment endpoints. Random fragmentation of the path is performed between MC steps, so that fixed system positions at previous fragment endpoints can be sampled during subsequent steps (Fig. 4.5C).

Table 4.1. Summary of notation employed to describe the PAMD integration scheme. The last three rows define parameters specific to the path sampling algorithm used.

dt	time-step for discretization of the sampled path
N_{path}	number of time-steps in the sampled path
N_{procs}	number of processors for parallel-in-time force evaluations
N_{force}	number of force evaluations per shifting event
N_{sample}	number of MC steps per shifting event
N_{shift}	number of time-steps shifted
L	total number of levels in the sampled path
l_{min}	finest sampled level of the path
l_{max}	coarsest sampled level of the path

At each MC step, a level is randomly selected between l_{min} and l_{max}, with $1 \leq l_{\text{min}} \leq l_{\text{max}} \leq L$. The calculations reported here employ $l_{\text{min}} > 1$ and $l_{\text{max}} < L$, such that not all levels are directly sampled. The choice of $l_{\text{max}} < L$ corresponds to excluding the direct sampling of levels associated with long path fragments, on the basis of negligible acceptance. The choice of $l_{\text{min}} > 1$ corresponds to excluding the direct sampling of levels associated with short path fragments, as these are trivially updated via the direct sampling of longer fragments at coarser levels. Despite these choices, the sampling remains ergodic due to the random fragmentation of the path that occurs between MC steps [38, 54].

Our choice of the multilevel S&S algorithm over other sampling algorithms offering gradient-based global path updates, based on an effective Hamiltonian or Langevin dynamics on the path space [47, 55–58], is motivated by at least three practical considerations specific to the context where PAMD applications are envisioned. The first pertains to the computational cost of each MC sampling step, which varies drastically across algorithms. While gradient-based algorithms, at least in their

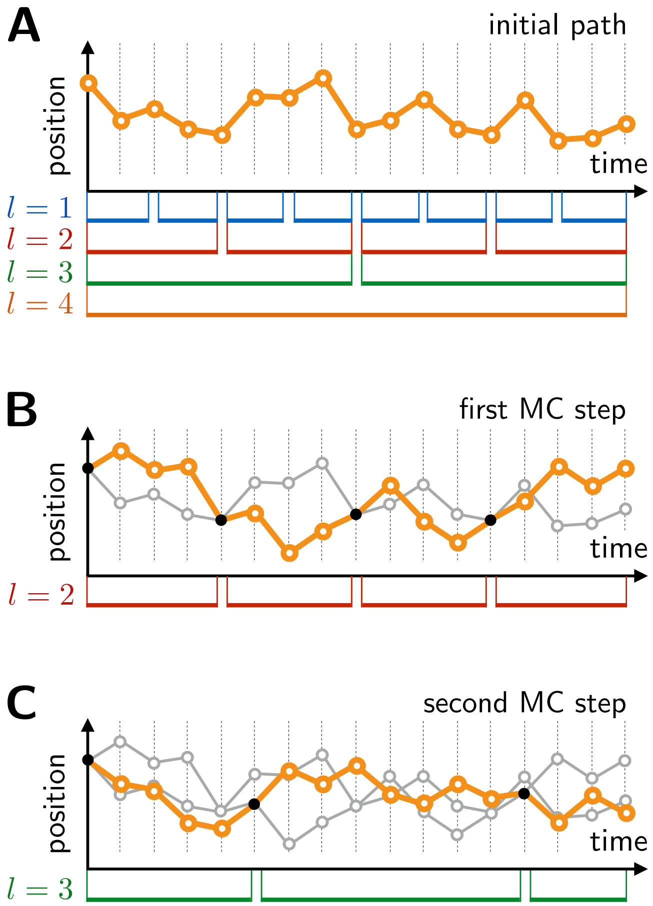

Figure 4.5. Illustration of the multilevel sliding and sampling (S&S) algorithm. **A** Multilevel representation of a path with $N_{\text{path}} = 16$ time-steps ($L = 4$ levels), shown for one Cartesian coordinate. An initial configuration of the path is shown in orange. **B** Update of the path configuration in **A** via a MC step at level $l = 2$. Fixed system positions along the path are represented with black dots, configurations of the path before the update in gray, and the configuration of the path after the update in orange. **C** MC step at level $l = 3$ following that shown in **B**. The new fragmentation of the path allows for updates of system positions that were held fixed in previous MC steps.

Metropolis-unadjusted variants, can provably exhibit faster mixing than gradient-free counterparts and exhibit robustness to arbitrary refinement of the MD time-step toward the continuous-time limit [55–58], they require the evaluation of higher derivatives of the physical potential that are often unavailable or considered too expensive in applications. On the other hand, gradient-free path sampling algorithms can be combined with a symmetrized choice of the discretized path-action functional [38, 45] to fully bypass all need for higher-order potential derivatives.

The second consideration is the requirement that the sampled path ensemble not exhibit spurious artifacts associated with the sampling method employed or the chosen discretization of the Onsager–Machlup functional. Observations have surfaced that gradient-based path sampling algorithms can produced unphysical path ensembles when combined with certain discretizations of the Onsager–Machlup functional [41, 59]; such artifacts are also avoided with gradient-free updates such as the celebrated Crank–Nicolson [60] (or Brownian tube [48]) method.

The third and final consideration is to avoid the slow MC mixing associated with global updates in long paths. The above-mentioned gradient-based and gradient-free methods are all formulated to induce updates of the full configuration of the sampled path in each MC step. The acceptance rate of such global updates can scale poorly with sampled path length at fixed time-step size [51, 61], and thus lead to slow equilibration within the long sampled-path lengths envisioned to be necessary for effective application of the PAMD algorithm. This scaling consideration also jeopardizes the performance of semi-local spectral Metropolis-within-Gibbs schemes that sequentially update the components of a normal-mode representation of the sampled path across its full length [62, 63].

The distribution of paths used to generate trial configurations for multilevel S&S in each application is chosen to maximize the statistical efficiency of the sampling (i.e., minimize N_{sample}) without requiring evaluation of the MD forces. For the harmonic oscillator, trials are drawn from the distribution of free particle paths. For simulations of the Lennard–Jones liquid, trials are drawn from the path distribution of a fluid of hard spheres with diameter σ_{HS}; this strategy reduces the number of force evaluations needed to obtain likely Lennard–Jones path configurations by excluding those with high interparticle overlap from the ensemble of trial paths. The likelihood of a path at the hard-sphere level is evaluated using an approximation of the pair propagator for diffusive hard spheres [64]. For both the harmonic oscillator and the Lennard–Jones applications, regeneration of the sampled path after shifting is performed with the same distribution used to generate trial configurations for the path sampling.

In total, each MC step involves a total number of N_{path} MD force evaluations. Since N_{force} is defined as the number of force evaluations per shifting event, and since N_{sample} is the number of MC steps per shifting event, we have

$$N_{\text{force}} = N_{\text{sample}} \cdot N_{\text{path}} \ . \tag{4.9}$$

As we seek to maximize the wall-clock speedup via parallelization of these independent force evaluations, we employ one processor per force evaluation, and thus

$$N_{\text{procs}} = N_{\text{path}} \ . \tag{4.10}$$

Additional parallelization within the force evaluation is of course possible, but is not considered in the current chapter. Thus, we insert Eqs. (4.9) and (4.10) into Eq. (4.8) to arrive at the following expression for the PAMD speedup:

$$\chi = \frac{N_{\text{shift}}}{N_{\text{sample}}} \cdot \frac{dt}{dt_{\text{E}}} \ . \tag{4.11}$$

See Tbl. 4.1 for a summary of terms.

4.4 Results

We now apply the PAMD algorithm to two model systems, considering time evolution according to Brownian dynamics (Eq. (4.1)), with $\beta = 1$ unless otherwise indicated and with $\gamma = 1$ in appropriately reduced units. Like the Euler algorithm (Eq. (4.2)), PAMD is a rigorous and formally exact way to integrate Brownian dynamics, yet the numerical accuracy of the trajectories depends on the parameters employed. In each application, we examine the relationship between the number of parallel processors employed and the speedup in the PAMD algorithm relative to the Euler algorithm (χ in Eq. (4.11)), subject to the requirement that the MD trajectories integrated using both PAMD and the Euler algorithm preserve well-defined measures of accuracy.

In the current chapter, we focus exclusively on wall-clock speedups achieved via parallelization of the MD integration in time, setting aside the separate and complementary issue of parallelizing the force evaluation at each time-step. All reported speedups for PAMD in the current study are theoretical; they are obtained from

Eq. (4.11) under the stated assumptions. The computational cost associated with each speedup can be obtained by multiplying the reported value with the number of parallel processors, N_{procs}, employed in each application.

4.4.1 Harmonic oscillator

Here, we consider the example of an overdamped harmonic oscillator, with potential $V(x) = \frac{1}{2}x^2$. Two measures of the accuracy of the integrated MD trajectories are considered. The first reports on the degree to which the trajectories sample the correct equilibrium distribution,

$$E_{\text{eq}} = \frac{\sqrt{\int_{\mathbb{R}} \mathrm{d}x \, |P(x) - P_{\text{s}}(x)|^2}}{Z} \tag{4.12}$$

where $P(x)$ is the exact Boltzmann distribution, Z is the associated partition function, and $P_{\text{s}}(x)$ is the equilibrium distribution of positions sampled by the numerical integration schemes. The second measure of error reports on the accuracy of the MD time evolution. Specifically, we consider the autocovariance function

$$C(t) = \langle x(t')x(t'+t) \rangle = \lim_{T \to \infty} \frac{1}{T} \int_0^T \mathrm{d}t' \, x(t')x(t'+t) \ , \tag{4.13}$$

which is a simple exponential function for the overdamped harmonic oscillator [49],

$$C(t) = \left\langle x^2 \right\rangle \exp(-\kappa t) \ , \tag{4.14}$$

where the angled brackets indicate Boltzmann averaging. The second measure of error is thus

$$E_{\text{dyn}} = \frac{|\kappa - \kappa_{\text{s}}|}{k} \tag{4.15}$$

where $\kappa = \gamma^{-1} = 1$ is the exact decay constant and κ_{s} is the decay constant obtained by fitting the exponential decay of the autocovariance from the numerically integrated MD trajectories. Specifically, κ_{s} is obtained by averaging over 100 independent trajectories of length 10^5 time units that are divided into 10^3 non-overlapping time series, for which the log-autocovariance is linearly fit in the range $t \in [0, 4]$. Simulation parameters for integration of the MD trajectories in this application are chosen to ensure that both measures of error remain below 3%.

Table 4.2 indicates that with a time-step of $dt_{\text{E}} = 0.025$, the Euler algorithm yields error values of $E_{\text{eq}} = 0.3\%$ and $E_{\text{dyn}} = 1.0\%$. Also shown in the table are parameters for three separate PAMD simulations that obtain speedups of $\chi = 16$, 128, and

1024 with respect to the Euler algorithm. For both the equilibrium distribution and the autocovariance function, Fig. 4.6 shows the comparison of the exact results and those obtained using PAMD with the aforementioned speedups. The accuracy of the PAMD trajectories is clearly preserved in all simulations, as indicated by the plotted results and the reported values of E_{eq} and E_{dyn}.

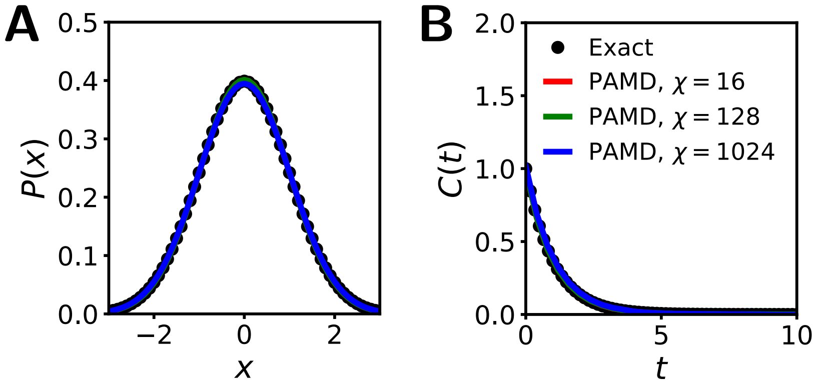

Figure 4.6. For the harmonic oscillator, comparison of PAMD results (colored lines) with exact results (black dots) for **A** the Boltzmann distribution $P(x)$ and **B** the autocovariance function $C(t)$. The PAMD results correspond to Simulation 1 (red), Simulation 2 (green), and Simulation 3 (blue) in Tbl. 4.2, which respectively achieve speedups of $\chi = 16$, 128, and 1024 relative to the Euler algorithm.

For all PAMD simulations reported in Tbl. 4.2, a significant fraction of the speedup comes from the 16-fold larger time-step that can be employed in the path-based scheme ($dt = 0.4$ vs. $dt_{\mathrm{E}} = 0.025$). The larger speedups achieved in Simulation 2 ($\chi = 128$) and Simulation 3 ($\chi = 1024$), in comparison to Simulation 1 ($\chi = 16$), arise from the larger ratios of N_{shift} to N_{sample} that are used in these simulations ($N_{\mathrm{shift}}/N_{\mathrm{sample}} = 8$ for Simulation 2 and $N_{\mathrm{shift}}/N_{\mathrm{sample}} = 64$ for Simulation 3) in comparison to Simulation 1 (for which $N_{\mathrm{shift}}/N_{\mathrm{sample}} = 1$). The higher frequency of shifting events associated with larger values of the ratio $N_{\mathrm{shift}}/N_{\mathrm{sample}}$ places greater demand on the efficiency of the path sampling, and a larger number of time-steps in the sampled path (N_{path}; hence, a larger number of parallel processors, N_{procs}, per Eq. (4.10)) is needed to allow path segments to undergo a sufficient number of MC steps before they are used to generate the marginal distribution for the MD trajectories, as discussed in connection with Figs. 4.3 and 4.4. Accordingly, the value of N_{path} in each of the simulations in Tbl. 4.2 is set as large as needed to ensure accurate MD integration at the corresponding value of N_{sample}, which is kept small ($N_{\mathrm{sample}} = 2$ in Simulation 1, and $N_{\mathrm{sample}} = 1$ in Simulation 2 and Simulation 3) to enhance the PAMD speedup.

Table 4.2. Summary of PAMD simulation parameters used for the harmonic oscillator application. Parameters L, l_{min}, and l_{max} are specific to the path sampling algorithm used, described in Sec. 4.3.

	Euler		
dt_{E}	0.025		
E_{eq}	0.3%		
E_{dyn}	1.0%		
	PAMD		
	Simulation 1	Simulation 2	Simulation 3
dt	0.4	0.4	0.4
N_{path}	16	256	4096
N_{sample}	2	1	1
N_{shift}	2	8	64
L	4	8	12
l_{min}	3	4	4
l_{max}	3	5	8
χ	16	128	1024
E_{eq}	0.3%	0.3%	0.5%
E_{dyn}	0.3%	1.1%	2.9%

4.4.2 Lennard–Jones liquid

Here, we apply PAMD to a model for a molecular liquid. The pairwise interaction between particles is described using the standard energy- and force-shifted Lennard–Jones potential [2],

$$U(r) = \begin{cases} u(r) - u(r_{\mathrm{c}}) - (r - r_{\mathrm{c}})\, u'(r_{\mathrm{c}}), & r \leq r_{\mathrm{c}} \\ 0, & r > r_{\mathrm{c}} \end{cases}, \tag{4.16}$$

where $u(r) = 4\epsilon\left\{(\sigma/r)^{12} - (\sigma/r)^{6}\right\}$; throughout, we take $\epsilon = 1$ and $\sigma = 1$. The system consists of 27 particles placed in a cubic box at reduced density $\rho\sigma^3 = 0.50$ and at constant reduced inverse temperature $\beta\epsilon = 0.74$. Simulations are performed with periodic boundary conditions at constant volume, and the cutoff distance r_{c} corresponds to half of the simulation box-length.

As described in Sec. 4.3, the reported simulations for the Lennard–Jones liquid employ trial configurations drawn from a distribution of paths for a fluid of hard spheres with diameter σ_{HS}. To prevent the path-sampling bias from affecting the accuracy of the integrated trajectories, we employ a hard-sphere schedule that varies as a function of the path-time τ, $\sigma_{\mathrm{HS}}(\tau)$; path configurations are then sampled in accordance with the path-time dependent potential

$$V(r;\sigma_{\mathrm{HS}}(\tau)) = U(r) + U_{\mathrm{HS}}(r;\sigma_{\mathrm{HS}}(\tau)) \;, \tag{4.17}$$

where $U(r)$ is defined in Eq. (4.16), and $U_{\mathrm{HS}}(r;\sigma_{\mathrm{HS}})$ is the hard-sphere potential

$$U_{\mathrm{HS}}(r;\sigma_{\mathrm{HS}}) = \begin{cases} +\infty, & r \leq \sigma_{\mathrm{HS}} \\ 0, & r > \sigma_{\mathrm{HS}} \end{cases} . \tag{4.18}$$

The schedule $\sigma_{\mathrm{HS}}(\tau)$ is chosen such that configurational volume is excluded at the nose of the path for enhanced sampling efficiency, and no volume is excluded at the tail of the path where the marginal distribution for MD integration is sampled (Fig. 4.7); in this way, path segments regenerated from the hard-sphere distribution are subsequently relaxed into the Lennard–Jones distribution as they shift from the nose to the tail of the path. A worthy direction for future work is to systematically optimize the schedule $\sigma_{\mathrm{HS}}(\tau)$, which in general will depend on the simulation parameters reported in Tbl. 4.1.

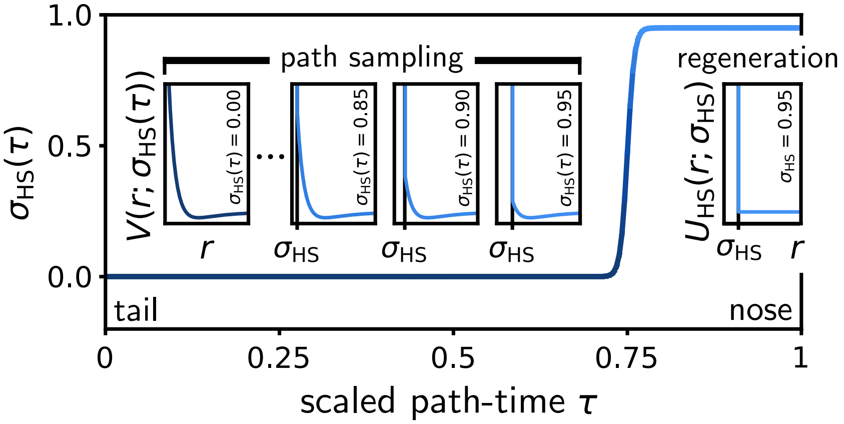

Figure 4.7. Hard-sphere schedule $\sigma_{\mathrm{HS}}(\tau)$, versus scaled path-time τ, employed in the reported Lennard–Jones simulations. Insets show slices of the path-time dependent potential $V(r;\sigma_{\mathrm{HS}}(\tau))$ employed to sample path configurations (Eq. (4.17)) at schedule values corresponding to biased ($\sigma_{\mathrm{HS}} \simeq 1$) and unbiased ($\sigma_{\mathrm{HS}} = 0$) sampling of the equilibrium path distribution, and the hard-sphere potential $U(r;\sigma_{\mathrm{HS}})$ used to regenerate path segments (Eq. (4.18)).

The accuracy of the integrated MD trajectories is evaluated in terms of the radial distribution function $g(r)$ and the self-diffusion coefficient D, using the respective error measures

$$E_{\mathrm{eq}} = \frac{\sqrt{\int_0^{r_{\mathrm{c}}} \mathrm{d}r \, |g(r) - g_{\mathrm{s}}(r)|^2}}{\int_0^{r_{\mathrm{c}}} \mathrm{d}r \, g(r)} \tag{4.19}$$

and

$$E_{\mathrm{dyn}} = \frac{|D - D_{\mathrm{s}}|}{D} \, , \tag{4.20}$$

where $g(r)$ and D are reference quantities obtained using the Euler algorithm with a small time-step (5×10^{-5} Lennard–Jones time units), and $g_{\mathrm{s}}(r)$ and D_{s} are obtained using PAMD and the Euler algorithm with larger time-steps. The diffusion coefficient is given by $D = \frac{1}{6} \lim_{t\to\infty} \frac{\mathrm{d}}{\mathrm{d}t} \langle R^2(t) \rangle$, where

$$\left\langle R^2(t) \right\rangle = \lim_{T\to\infty} \int_0^T \mathrm{d}t' \, \frac{1}{N} \sum_{i=1}^{N} |\mathbf{r}_i(t'+t) - \mathbf{r}_i(t')|^2 \tag{4.21}$$

is the mean-square displacement, $\mathbf{r}_i$ the position of the ith particle, and N the number of particles [1]. $\langle R^2(t) \rangle$ is obtained by averaging over 100 independent trajectories that are divided into 100 non-overlapping time series of length 1, in Lennard–Jones time units, and a linear fit is performed in the range $t \in [0.2, 1]$ to evaluate D_{s} for the PAMD and Euler simulations.

Table 4.3 indicates that at a time-step of $dt_{\mathrm{E}} = 2.5 \times 10^{-4}$, the Euler algorithm yields error values of $E_{\mathrm{eq}} = 1.4\%$ and $E_{\mathrm{dyn}} = 0.1\%$; larger time-steps were found to lead to unstable Euler trajectories. Also shown in Tbl. 4.3 are two PAMD simulations that lead to 16-fold ($\chi = 16$; Simulation 1) and 128-fold ($\chi = 128$; Simulation 2) reductions of the wall-clock time required to generate equivalently accurate MD trajectories for the Lennard–Jones liquid via the Euler algorithm, using simulation parameters that keep error values below 5%. The radial distribution functions and mean-square displacements obtained from these two simulations are plotted with the corresponding reference quantities in Figs. 4.8A and B. Excellent agreement between the PAMD and reference quantities is evident in the plots and from the values of E_{eq} and E_{dyn} reported in Tbl. 4.3.

As in the harmonic oscillator application, the speedups reported in Tbl. 4.3 for the Lennard–Jones liquid are partially enabled by the use of a larger time-step in PAMD ($dt = 5 \times 10^{-4}$) than is possible for stable numerical integration via the Euler method ($dt_{\mathrm{E}} = 2.5 \times 10^{-4}$). The remaining speedup in both simulations comes from using shift lengths that integrate $N_{\mathrm{shift}} = 8$ (Simulation 1) and $N_{\mathrm{shift}} = 64$ (Simulation 2) time-steps of MD trajectory at a rate of $N_{\mathrm{sample}} = 1$ MC steps per shifting event. Accurate integration at these speedups requires efficient sampling of path modes that are commensurate with the shifting timescale ($8\, dt$ in Simulation 1 and $64\, dt$ in Simulation 2); accordingly, long paths ($N_{\mathrm{path}} = 128$ in Simulation 1 and $N_{\mathrm{path}} = 4096$ in Simulation 2) are employed in both simulations so that segments can undergo a

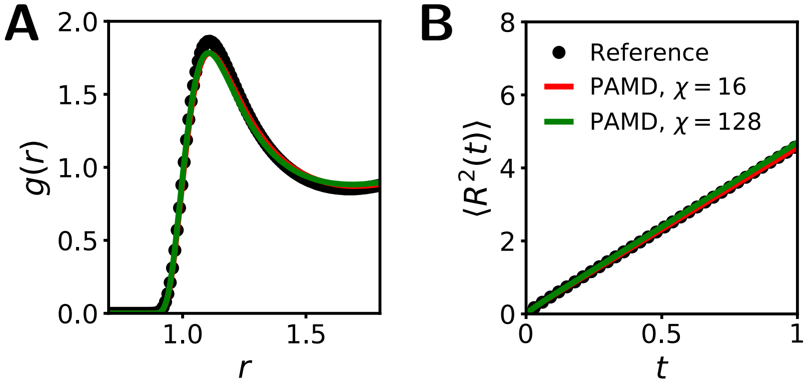

Figure 4.8. For the Lennard–Jones liquid at $\beta\epsilon = 0.74$ and $\rho\sigma^3 = 0.50$, comparison of PAMD results (colored lines) with numerically exact results (black dots) for **A** the radial distribution function and **B** the mean-square displacement. The PAMD results correspond to Simulation 1 (red) and Simulation 2 (green) in Tbl. 4.3, which respectively achieve speedups of $\chi = 16$ and 128 relative to the Euler algorithm.

sufficient number of MC steps before they are used to generate the respective MD trajectories (Fig. 4.4).

4.5 Conclusions

The field of MD simulation faces important challenges in harnessing massively parallel computer architectures. Although successful parallelization of the the force evaluation can be expected as the system size grows (i.e., weak scaling), there exists a much more difficult challenge of employing ever-larger numbers of parallel processors to accelerate the simulation of systems of a fixed size (i.e., strong scaling). Remarkable success has been achieved in this vein [8–13], but fundamental limitations are inevitable.

The work comprising this chapter suggests that parallelization in the dimension of time via path integrals offers a promising avenue for future progress in terms of strong scaling. We introduce the PAMD approach, which achieves significant speedups over conventional algorithms for stochastic dynamics by parallelizing evaluation of the MD forces along trajectory segments with respect to time. The method strives to be applicable in greater generality than alternative approaches to long-timescale simulation of stochastic trajectories, such as milestoning [65], Markov state modeling [66], and trajectory splicing [24], by enabling the direct simulation of long-timescale processes without assuming separation of timescales, and without re-

Table 4.3. Summary of PAMD simulation parameters used for the application to the Lennard–Jones liquid. L, l_{min}, and l_{max} are parameters specific to the path sampling algorithm used, described in Sec. 4.3.

	Euler	
dt_{E}	2.5×10^{-4}	
E_{eq}	1.4%	
E_{dyn}	0.1%	
	PAMD	
	Simulation 1	Simulation 2
dt	5×10^{-4}	5×10^{-4}
N_{path}	128	4096
N_{sample}	1	1
N_{shift}	8	64
L	7	12
l_{min}	3	4
l_{max}	5	7
χ	16	128
E_{eq}	3.5%	3.7%
E_{dyn}	1.0%	1.9%

quiring the specification of a suitable configuration-space tiling or a set of collective variables. Proof-of-principle applications to overdamped systems show that PAMD can accelerate trajectory integration by several orders of magnitude with respect to the conventional Euler scheme for Brownian dynamics, and even greater speedups are possible with the use of larger numbers of parallel processors and enhancement of the MC path sampling efficiency.

Looking forward, the PAMD approach will likely require additional methodological developments to become viable for large-scale simulations of complex systems. Central to this effort will be the refinement of path sampling methodologies that lead to the reduction in the number of parallel processors that are needed for a given amount of speedup with the method. Our application to the Lennard–Jones liquid shows that preconditioned sampling of paths from an approximate dynamical model [67], constructed *a priori* based on known features of the dynamics or generated on-the-fly via adaptive inference techniques for stochastic process approximation [68–70], provides an avenue for efficiency improvements in more general applications. Regardless, we feel that the natural parallelization of path-integral formulations, combined with the increasing availability of massively parallel computer resources, should motivate increased attention to the opportunities of parallelizing molecular dynamics simulation in time.

Supplementary Material

Supplementary movies M1 and M2, respectively, show execution of the PAMD integration scheme as it accumulates position histograms during simulations of the Brownian harmonic oscillator, for the cases $N_{\mathrm{sample}} = 1$ (blue curve) and $N_{\mathrm{sample}} = 4$ (red curve) in Fig. 4.3E. These cases illustrate distinct operational regimes of the PAMD algorithm, respectively corresponding to inaccurate and accurate stochastic MD trajectories as a function number of path sampling steps per PAMD iteration. The movies are available from CaltechDATA at `DOI: 10.22002/D1.20111`.

Acknowledgments. We are grateful for stimulating discussions with Matthew G. Welborn, Eric Vanden-Eijnden, and Gavin E. Crooks. The work comprising this chapter was supported in part by the Department of Energy under Award No. DE-FOA-0001912 and the Office of Naval Research under Award No. N00014-10-1-0884. This research also used resources of the National Energy Research Scientific Computing Center (NERSC), a DOE Office of Science User Facility supported by the DOE Office of Science under contract DE-AC02-05CH11231.

References Cited in Part II

[1] D. Frenkel and B. Smit, *Understanding molecular simulation: From algorithms to applications* (Academic Press, 2002).

[2] M. P. Allen and D. J. Tildesley, *Computer simulation of liquids* (Clarendon Press, 2017).

[3] J. D. Durrant and J. A. McCammon, "Molecular dynamics simulations and drug discovery", BMC Biology **9**, 71 (2011).

[4] R. O. Dror, R. M. Dirks, J. P. Grossman, H. Xu, and D. E. Shaw, "Biomolecular simulation: A computational microscope for molecular biology", Annual Review of Biophysics **41**, 429 (2012).

[5] A. F. Voter, F. Montalenti, and T. C. Germann, "Extending the time scale in atomistic simulation of materials", Annual Review of Materials Research **32**, 321 (2002).

[6] D. Keyes, in *High performance computing - HiPC 2007*, edited by S. Aluru, M. Parashar, R. Badrinath, and V. K. Prasanna (Springer Berlin Heidelberg, 2007), pp. 2–3.

[7] J. Dongarra, P. Beckman, T. Moore, P. Aerts, G. Aloisio, J.-C. Andre, D. Barkai, J.-Y. Berthou, T. Boku, B. B, F. Cappello, B. Chapman, X. Chi, A. Choudhary, S. Dosanjh, T. Dunning, S. Fiore, A. Geist, B. Gropp, B. Harrison, M. Hereld, M. Heroux, A. Hoisie, K. Hotta, Z. Jin, Y. Ishikawa, F. Johnson, S. Kale, R. Kenway, D. Keyes, B. Kramer, J. Labarta, A. Lichnewsky, T. Lippert, B. Lucas, B. Maccabe, S. Matsuoka, P. Messina, P. Michielse, M. B., M. S. Mueller, W. E. Nagel, H. Nakashima, M. E. Papka, D. Reed, M. Sato, E. Seidel, J. Shalf, D. Skinner, M. Snir, T. Sterling, R. Stevens, F. Streitz, B. Sugar, S. Sumimoto, W. Tang, J. Taylor, R. Thakur, A. Trefethen, M. Valero, A. van der Steen, J. Vetter, P. Williams, R. Wisniewski, and K. Yelick, "The International Exascale Software Project roadmap", The International Journal of High Performance Computing Applications **25**, 3 (2011).

[8] S. Plimpton, "Fast parallel algorithms for short-range molecular-dynamics", Journal of Computational Physics **117**, 1 (1995).

[9] J. C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kalé, and K. Schulten, "Scalable molecular dynamics with NAMD", Journal of Computational Chemistry **26**, 1781 (2005).

[10] D. E. Shaw, R. O. Dror, J. K. Salmon, J. P. Grossman, K. M. Mackenzie, J. A. Bank, C. Young, M. M. Deneroff, B. Batson, K. J. Bowers, E. Chow, M. P. Eastwood, D. J. Ierardi, J. L. Klepeis, J. S. Kuskin, R. H. Larson, K. Lindorff-Larsen, P. Maragakis, M. A. Moraes, S. Piana, Y. Shan, and B. Towles, "Millisecond-scale molecular dynamics simulations on Anton", in Proceedings of

the conference on high performance computing networking, storage and analysis, SC ’09 (2009), pp. 1–11.

[11] R. Salomon-Ferrer, A. W. Götz, D. Poole, S. Le Grand, and R. C. Walker, “Routine microsecond molecular dynamics simulations with AMBER on GPUs. 2. Explicit solvent particle mesh Ewald”, Journal of Chemical Theory and Computation **9**, 3878 (2013).

[12] S. Páll, M. J. Abraham, C. Kutzner, B. Hess, and E. Lindahl, in *Solving software challenges for exascale*, edited by S. Markidis and E. Laure (Springer International Publishing, 2015), pp. 3–27.

[13] J. P. Grossman, B. Towles, B. Greskamp, and D. E. Shaw, “Filtering, reductions and synchronization in the Anton 2 network”, Parallel and Distributed Processing Symposium, 860 (2015).

[14] J.-L. Lions, Y. Maday, and G. Turinici, “A parareal discretization in time of PDE’s”, Comptes Rendus de l’Académie des Sciences - Series I - Mathematics **332**, 661 (2001).

[15] C. Farhat and M. Chandesris, “Time-decomposed parallel time-integrators: Theory and feasibility studies for fluid, structure, and fluid-structure applications”, International Journal for Numerical Methods in Engineering **58**, 1397 (2003).

[16] I. D. Garrido, B. Lee, G. E. Fladmark, and M. S. Espedal, “Convergent iterative schemes for time parallelization”, Mathematics of Computation **75**, 1403 (2006).

[17] M. Emmett and M. L. Minion, “Toward an efficient parallel in time method for partial differential equations”, Communications in Applied Mathematics and Computational Science **7**, 105 (2012).

[18] L. Baffico, S. Bernard, Y. Maday, G. Turinici, and G. Zérah, “Parallel-in-time molecular-dynamics simulations”, Physical Review E **66**, 057701 (2002).

[19] Y. Yanan, A. Srinivasan, and N. Chandra, “Scalable time-parallelization of molecular dynamics simulations in nano mechanics”, in Proceedings of the 2006 international conference on parallel processing, ICPP ’06 (2006), pp. 119–126.

[20] R. Speck, D. Ruprecht, R. Krause, M. Emmett, M. Minion, M. Winkel, and P. Gibbon, “A massively space-time parallel N-body solver”, in Proceedings of the international conference on high performance computing, networking, storage and analysis, SC ’12 (2012), pp. 92-1–92-11.

[21] E. J. Bylaska, J. Q. Weare, and J. H. Weare, “Extending molecular simulation time scales: Parallel in time integrations for high-level quantum chemistry and complex force representations”, The Journal of Chemical Physics **139**, 074114 (2013).

[22] A. L. Blumers, Z. Li, and G. E. Karniadakis, “Supervised parallel-in-time algorithm for long-time Lagrangian simulations of stochastic dynamics: Application to hydrodynamics”, Journal of Computational Physics **393**, 214 (2019).

[23] F. Legoll, T. Lelièvre, and U. Sharma, “An adaptive Parareal algorithm: Application to the simulation of molecular dynamics trajectories”, SIAM Journal on Scientific Computing **44**, B146 (2022).

[24] D. Perez, E. D. Cubuk, A. Waterland, E. Kaxiras, and A. F. Voter, “Long-time dynamics through parallel trajectory splicing”, Journal of Chemical Theory and Computation **12**, 18 (2016).

[25] A. Brünger, C. L. Brooks, and M. Karplus, “Stochastic boundary conditions for molecular dynamics simulations of ST2 water”, Chemical Physics Letters **105**, 495 (1984).

[26] A. C. Brańka and D. M. Heyes, “Algorithms for Brownian dynamics simulation”, Physical Review E **58**, 2611 (1998).

[27] G. Bussi and M. Parrinello, “Accurate sampling using Langevin dynamics”, Physical Review E **75**, 056707 (2007).

[28] N. Bou-Rabee, “Time integrators for molecular dynamics”, Entropy **16**, 138 (2014).

[29] D. A. Sivak, J. D. Chodera, and G. E. Crooks, “Time step rescaling recovers continuous-time dynamical properties for discrete-time Langevin integration of nonequilibrium systems”, Journal of Physical Chemistry B **118**, 6466 (2014).

[30] H. Risken and T. Frank, *The Fokker–Planck equation: Methods of solution and applications* (Springer Berlin Heidelberg, 1996).

[31] L. Onsager and S. Machlup, “Fluctuations and irreversible processes”, Physical Review **91**, 1505 (1953).

[32] K. Yasue, “A simple derivation of the Onsager–Machlup formula for one-dimensional nonlinear diffusion process”, Journal of Mathematical Physics **19**, 1671 (1978).

[33] L. R. Pratt, “A statistical method for identifying transition states in high dimensional problems”, The Journal of Chemical Physics **85**, 5045 (1986).

[34] R. Olender and R. Elber, “Calculation of classical trajectories with a very large time step: Formalism and numerical examples”, The Journal of Chemical Physics **105**, 9299 (1996).

[35] D. M. Zuckerman and T. B. Woolf, “Efficient dynamic importance sampling of rare events in one dimension”, Physical Review E **63**, 016702 (2000).

[36] D. Passerone and M. Parrinello, “Action-derived molecular dynamics in the study of rare events”, Physical Review Letters **87**, 108302 (2001).

[37] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, "Transition path sampling: Throwing ropes over rough mountain passes, in the dark", Annual Reviews of Physical Chemistry **53**, 291 (2002).

[38] T. F. Miller III and C. Predescu, "Sampling diffusive transition paths", The Journal of Chemical Physics **126**, 144102 (2007).

[39] J. MacFadyen, J. Wereszczynski, and I. Andricioaei, "Directionally negative friction: A method for enhanced sampling of rare event kinetics", The Journal of Chemical Physics **128**, 114112 (2008).

[40] W. E and E. Vanden-Eijnden, "Transition-path theory and path-finding algorithms for the study of rare events", Annual Review of Physical Chemistry **61**, 391 (2010).

[41] F. J. Pinski and A. M. Stuart, "Transition paths in molecules at finite temperature", The Journal of Chemical Physics **132**, 184104 (2010).

[42] H. Orland, "Generating transition paths by Langevin bridges", The Journal of Chemical Physics **134**, 174114 (2011).

[43] S. N. Majumdar and H. Orland, "Effective Langevin equations for constrained stochastic processes", Journal of Statistical Mechanics: Theory and Experiment **2015**, P06039 (2015).

[44] P. D. Drummond, "Forward, backward, and weighted stochastic bridges", Physical Review E **96**, 042123 (2017).

[45] J. Lee, I.-H. Lee, I. Joung, L. Jooyoung, and B. R. Brooks, "Finding multiple reaction pathways via global optimization of action", Nature Communications **8**, 15443 (2017).

[46] D. M. Ceperley, "Path integrals in the theory of condensed helium", Reviews of Modern Physics **67**, 279 (1995).

[47] A. M. Stuart, J. Voss, and P. Wilberg, "Conditional path sampling of SDEs and the Langevin MCMC method", Communications in Mathematical Sciences **2**, 685 (2004).

[48] G. Stoltz, "Path sampling with stochastic dynamics: Some new algorithms", Journal of Computational Physics **225**, 491 (2007).

[49] C. W. Gardiner, *Stochastic methods: A handbook for the natural and social sciences*, Springer Series in Synergetics (Springer Berlin Heidelberg, 2009).

[50] C. Dellago, P. G. Bolhuis, and P. L. Geissler, in *Advances in chemical physics* (John Wiley & Sons, Inc., 2003), pp. 1–78.

[51] T. R. Gingrich and P. L. Geissler, "Preserving correlations between trajectories for efficient path sampling", The Journal of Chemical Physics **142**, 234104 (2015).

[52] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines", The Journal of Chemical Physics **21**, 1087 (1953).

[53] W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications", Biometrika **57**, 97 (1970).

[54] C. H. Mak and S. Zakharov, "A multigrid algorithm for sampling imaginary-time paths in quantum Monte Carlo simulations", Journal of Physical Chemistry B **108**, 6760 (2004).

[55] A. Beskos, G. Roberts, A. Stuart, and J. Voss, "MCMC methods for diffusion bridges", Stochastics and Dynamics **8**, 319 (2008).

[56] A. Beskos, F. J. Pinski, J. M. Sanz-Serna, and A. M. Stuart, "Hybrid Monte Carlo on Hilbert spaces", Stochastic Processes and their Applications **121**, 2201 (2011).

[57] M. Ottobre, N. S. Pillai, F. J. Pinski, and A. M. Stuart, "A function space HMC algorithm with second order Langevin diffusion limit", Bernoulli **22**, 60 (2016).

[58] A. Beskos, M. Girolami, S. Lan, P. E. Farrell, and A. M. Stuart, "Geometric MCMC for infinite-dimensional inverse problems", Journal of Computational Physics **335**, 327 (2017).

[59] P. J. Malsom and F. J. Pinski, "Role of Ito's lemma in sampling pinned diffusion paths in the continuous-time limit", Physical Review E **94**, 042131 (2016).

[60] M. Hairer, A. M. Stuart, J. Voss, and P. Wiberg, "Analysis of SPDEs arising in path sampling. Part I: The Gaussian case", Communications in Mathematical Sciences **3**, 587 (2005).

[61] A. B. Adib, "Stochastic actions for diffusive dynamics: Reweighting, sampling, and minimization", The Journal of Physical Chemistry B **112**, 5910 (2008).

[62] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White, "MCMC methods for functions: Modifying old algorithms to make them faster", Statistical Science **28**, 424 (2013).

[63] N. Bou-Rabee, "SPECTRWM: Spectral Random Walk Method for the Numerical Solution of Stochastic Partial Differential Equations", SIAM Review **60**, 386 (2018).

[64] H. Behringer and R. Eichhorn, "Brownian dynamics simulations with hard-body interactions: Spherical particles", The Journal of Chemical Physics **137**, 164108 (2012).

[65] G. Grazioli and I. Andricioaei, “Advances in milestoning. II. Calculating time-correlation functions from milestoning using stochastic path integrals”, The Journal of Chemical Physics **149**, 084104 (2018).

[66] L.-T. Da, F. K. Sheong, D.-A. Silva, and X. Huang, “Application of Markov state models to simulate long timescale dynamics of biological macromolecules”, in *Protein conformational dynamics*, edited by K.-L. Han, X. Zhang, and M.-J. Yang (Springer International Publishing, 2014), pp. 29–66.

[67] J. O. B. Tempkin, B. Qi, M. G. Saunders, B. Roux, A. R. Dinner, and J. Weare, “Using multiscale preconditioning to accelerate the convergence of iterative molecular calculations”, The Journal of Chemical Physics **140**, 184114 (2014).

[68] M. Opper, “Variational inference for stochastic differential equations”, Annalen der Physik **531**, 1800233 (2019).

[69] V. Harmandaris, E. Kalligiannaki, M. Katsoulakis, and P. Plecháč, “Path-space variational inference for non-equilibrium coarse-grained systems”, Journal of Computational Physics **314**, 355 (2016).

[70] S.-T. Tsai, E.-J. Kuo, and P. Tiwary, “Learning molecular dynamics with simple language model built upon long short-term memory neural network”, Nature Communications **11**, 5115 (2020).