

The inefficiency of re-weighted sampling and the curse of system size in high order path integration

Michele Ceriotti¹, Guy A. R. Brain¹, Oliver Riordan², and David E. Manolopoulos¹

¹*Physical and Theoretical Chemistry Laboratory, South Parks Road, Oxford OX1 3QZ, UK* ²*Mathematical Institute, 24-29 St Giles', Oxford OX1 3LB, UK*

Computing averages over a target probability density by statistical re-weighting of a set of samples with a different distribution is a strategy which is commonly adopted in fields as diverse as atomistic simulation and finance. Here we present a very general analysis of the accuracy and efficiency of this approach, highlighting some of its weaknesses. We then give an example of how our results can be used, specifically to assess the feasibility of high-order path integral methods. We demonstrate that the most promising of these techniques – which is based on re-weighted sampling – is bound to fail as the size of the system is increased, because of the exponential growth of the statistical uncertainty in the re-weighted average.

Key words: statistical sampling, re-weighting, high-order path integrals.

1. Introduction

Averaging is to the theoretical and computational sciences what measuring is to empirical science, and it is hard to imagine a problem with more general implications. Formally, the problem of computing an expectation value for a quantity $a(x)$ over a given probability distribution $p(x)$ can be stated as

$$\langle a \rangle_p = \int a(x)p(x)dx. \quad (1.1)$$

Very often, it turns out that x spans a very high dimensional space, which makes it impossible to evaluate the integral on a grid of points. Instead, importance sampling algorithms [1–3] are used which generate a set of points $x_i^{(p)}$ distributed in accordance with the target probability p . Then, the expectation value can be computed as an average over that set of points,

$$\langle a \rangle_p \approx \bar{a}_n^{(p)} = n^{-1} \sum_{i=1}^n a(x_i^{(p)}). \quad (1.2)$$

The error in the estimate (1.2) decreases with $n^{-1/2}$, *provided that the sample points are uncorrelated*. In a number of circumstances obtaining uncorrelated samples from p is an exceedingly difficult problem, either because generating an individual point is computationally demanding, or because the sampling algorithm produces strongly correlated points, so that each one contributes very little to the reduction of the statistical error.

In these cases one is led to explore the possibility of performing re-weighted sampling, i.e., generating points according to a different distribution $p_0(x)$, which is easier to sample efficiently, and then computing the average relative to p using a correction weight $w(x) = p(x)/p_0(x)$ involving the ratio of the probability densities:

$$\langle a \rangle_p = \frac{\langle aw \rangle_{p_0}}{\langle w \rangle_{p_0}} \approx \bar{a}_n = \frac{\sum_{i=1}^n a(x_i^{(p_0)}) w(x_i^{(p_0)})}{\sum_{i=1}^n w(x_i^{(p_0)})}. \quad (1.3)$$

Many methods have been developed which are fundamentally based on the re-weighted average (1.3), which has found applications in physical chemistry [4–6], theoretical physics [7, 8], economics and statistics [2, 9]. Section 2 will be devoted to a thorough analysis of the statistical properties of re-weighted sampling, and we will discuss under quite general assumptions the conditions required to apply this technique successfully.

Our analysis allows one to assess the efficacy of re-weighting in each of the contexts in which it is applied. One instance of an important problem which has recently been tackled by re-weighted sampling concerns the use of high-order path integration techniques in computer simulations which allow for the quantum nature of atomic motion. Conventional path integral (PI) methods [10–12] involve a significant overhead compared to a purely classical treatment, and a considerable effort has been devoted to the quest for less computationally demanding approaches. High order discretizations of the path have been available for some time [13–15], but it has only recently become apparent that one can circumvent the calculation of the bothersome second derivatives of the physical potential that arise in these discretizations by using re-weighted sampling [16, 17]. If successful, this approach would make high order path integration generally applicable and very attractive for many applications.

To provide an example of the utility of the results established in Section 2, we shall therefore proceed in Section 3 to focus on the role of re-weighted sampling in high order PI simulations. We shall show in particular that this form of path integration is eventually bound to fail because of the reduced statistical efficiency introduced by the re-weighting. Simulations of small clusters and/or mildly quantum mechanical problems constitute a niche in which this approach might be beneficial [17], but its performance is bound to degrade for larger systems.

2. Statistics of re-weighting

A key issue which is often overlooked is that averaging according to (1.3) will have a lower statistical efficiency than sampling p directly as in (1.2), i.e., for the same number of *uncorrelated* sample points, the statistical error in the re-weighted average will often be larger. Working out a simple, analytical expression for this drop in sampling efficiency would be extremely useful, as it would allow one to make an informed choice as to whether the computational gain of sampling p_0 is overshadowed by the concomitant statistical inefficiency.

In order to evaluate such an estimate, let us first simplify the notation of Eq. (1.3) by labelling a_i the i -th sample for the observable a , and w_i the ratio of the probability densities $p(x_i)/p_0(x_i)$. Averages with respect to the target distribution p will be written

as $\langle \cdot \rangle_p$, while averages relative to p_0 will be written as $\langle \cdot \rangle$. With this notation,

$$\bar{a}_n = \sum_{i=1}^n a_i w_i / \sum_{i=1}^n w_i, \quad \text{and} \quad \lim_{n \rightarrow \infty} \bar{a}_n = \frac{\langle aw \rangle}{\langle w \rangle} = \langle a \rangle_p. \quad (2.1)$$

Our objective is to unravel the statistical properties of the weighted average (2.1) for a finite number n of sample points, and to compare them with those resulting from sampling the target distribution directly. We will assume that different samples are completely uncorrelated, since different sampling distributions and algorithms can be compared on the basis of the computational cost of generating a new uncorrelated sample point.

It is shown in the appendix that an asymptotic expression for the expectation value and the variance of \bar{a}_n can be obtained under very weak assumptions about the joint probability distribution

$$p_0(a, w) = \int \delta(a - a(x)) \delta(w - w(x)) p_0(x) dx$$

of a and w . To first order in n^{-1} ,

$$\langle \bar{a}_n \rangle \approx \frac{\langle aw \rangle}{\langle w \rangle} + \frac{\langle aw \rangle \langle w^2 \rangle - \langle w \rangle \langle aw^2 \rangle}{n \langle w \rangle^3}, \quad (2.2)$$

$$\sigma^2(\bar{a}_n) \approx \frac{\langle a^2 w^2 \rangle \langle w \rangle^2 - 2 \langle aw \rangle \langle aw^2 \rangle \langle w \rangle + \langle aw \rangle^2 \langle w^2 \rangle}{n \langle w \rangle^4}. \quad (2.3)$$

Eq. (2.2) shows that for any finite number of samples the re-weighted average is a biased estimator of $\langle a \rangle_p$, i.e., that $\langle \bar{a}_n \rangle$ is affected by a systematic error which decreases with n^{-1} , while Eq. (2.3) shows that the statistical error in the re-weighted mean decreases asymptotically as $n^{-1/2}$.

In order to perform a comparison between the weighted and the un-weighted sampling strategies, it is necessary to evaluate the coefficients which enter Eqs. (2.2) and (2.3). To do this we introduce an additional assumption, namely that a and $h = -\ln w$ are correlated Gaussian variates, with means $\langle a \rangle$ and 0, variances $\sigma^2(a)$ and $\sigma^2(h)$, and covariance $\langle ah \rangle$, all with respect to p_0 . We introduce h because in many cases – certainly the vast majority of those occurring in the physical sciences – the probability density is proportional to the exponential of an extensive state function (e.g. a difference Hamiltonian). It is not difficult to justify the assumption of Gaussian statistics, since the observable and the logarithm of the weight are often computed by summing a large number of fluctuating contributions, and will therefore have nearly Gaussian statistics by virtue of the central limit theorem. In any case the objective here is simply to discuss the qualitative features of re-weighted sampling, and the Gaussian limit provides a convenient framework in which to do this.

Under the assumption of a Gaussian joint probability distribution for (a, h) , all expectation values of the form $\langle a^p w^q \rangle$ can be computed analytically in terms of the parameters of the distribution. Eqs. (2.2) and (2.3) then simplify to

$$\langle \bar{a}_n \rangle \approx \langle a \rangle - \langle ah \rangle + \langle ah \rangle e^{\sigma^2(h)}/n, \quad (2.4)$$

$$\sigma^2(\bar{a}_n) \approx \left(\sigma^2(a) + \langle ah \rangle^2 \right) e^{\sigma^2(h)}/n. \quad (2.5)$$

These equations convey in a more transparent way the message that is inherent in their more general counterparts. The systematic bias depends on the cross-correlation between the observable a and $h = -\ln w$, and both the systematic and statistical error grow exponentially with $\sigma^2(h)$. Whenever the fluctuations of h are smaller than 1, there will be little difference between the efficiency of sampling directly based on the target distribution and that of the re-weighted method. However, as soon as $\sigma(h)$ becomes larger than one, both the systematic bias and the statistical error of the re-weighted approach will require a much larger number of samples for convergence to within a given threshold.




Figure 1. Numerical results for the mean and standard deviation in the mean for n independent samples obtained from re-weighting a Gaussian-distributed observable with zero mean and variance $\sigma^2(a) = 1$, based on a difference Hamiltonian h with variance $\sigma^2(h)$, and cross-correlation $\langle ah \rangle$ with a . The full lines correspond to the predictions of Eqs. (2.4) and (2.5).

In Figure 1 we present the results of some numerical experiments in which we generated correlated Gaussian variables a and h , and computed the weighted average of a according to (1.3). The results confirm the asymptotic nature of our estimates, which describe the behaviour of $\langle \bar{a}_n \rangle$ and $\sigma^2(\bar{a}_n)$ in the large- n regime. The small- n limit can instead be pin-pointed very easily by considering the $n = 1$ case, in which the average converges to the un-weighted value $\langle a \rangle$, and the variance is just $\sigma^2(a)$.

While it is common knowledge that the fluctuations of the “difference Hamiltonian” $h = -\ln w$ must be small in order to apply re-weighting successfully, we believe it will come as a surprise to many that fluctuations of the order of a few units in h will decrease the sampling efficiency by orders of magnitude. The message here is that in re-weighted simulations the joint probability distribution of a and h should be monitored carefully. The value of $\exp[\sigma^2(h)]$ provides a direct indication of the statistical inefficiency due to re-weighting, and the cross-correlation $\langle ah \rangle$ provides an equally direct indication of the systematic bias in $\langle \bar{a}_n \rangle$. Both of these are straightforward to converge, as they involve un-weighted averages.

While Eqs. (2.4) and (2.5) cannot be used quantitatively in the case of non-Gaussian statistics, they provide useful guidelines to judge the quality of sampling in a re-weighted calculation. Whenever the fluctuations of h are large and the weighted average is much closer to $\langle a \rangle$ than to $\langle a \rangle - \langle ah \rangle$, it is very likely that the numerical results obtained from (1.3) will be meaningless. In order to provide an illustration of the utility of these estimates, we will now show that they can be used to predict the situations in which high-order PI schemes based on re-weighted sampling can be used successfully, and those in which they cannot.

3. High order path integration

(a) Suzuki-Chin path integration

Path integral methods aim at approximating the quantum mechanical partition function, $Z = \text{Tr}e^{-\beta H}$, by factoring it in such a way that the resulting expression can be evaluated within a classical framework. The simplest such factorization makes use of the well-known Trotter product formula [18], and results in an expression which is equivalent to a purely classical partition function in an extended phase space consisting of many replicas (beads) of the system, connected by harmonic springs to form a closed path (necklace) [10–12, 19]. The overall Trotter Hamiltonian reads

$$H^{(T)} = \sum_{j=1}^P \frac{p_j^2}{2m} + \frac{m}{2} \left(\frac{\beta \hbar}{P} \right)^2 (q_j - q_{j+1})^2 + V(q_j), \quad (3.1)$$

where $\{q_j\}$ and $\{p_j\}$ are the coordinates and momenta of the different beads, and $V(q)$ is the physical potential. The classical partition function computed with this Hamiltonian at P times the physical temperature converges to the fully quantum mechanical partition function Z , with an error that is second-order in β/P . The number of beads necessary for convergence is in general a small multiple of $\beta \hbar \omega_{max}$, where ω_{max} is the fastest physical vibration in the system. The computational cost therefore becomes very large for low temperatures and/or stiff vibrational modes, and considerable effort has been devoted to the quest for faster-converging factorizations in order to reduce this cost.

It has been shown [20] that factorizations with an error that is of an order higher than two in β/P must contain composite operators which have the form of nested commutators between the kinetic energy operator \hat{T} and the potential \hat{V} . The complexity of the nested commutators increases with the order of the factorization and is only really manageable up to fourth-order. Among the different fourth-order factorizations developed so far [13–16, 21], the one proposed by Suzuki [14] and simplified by Chin [15] has better convergence properties than most others. It also only contains the double commutator $[\hat{V}, [\hat{T}, \hat{V}]]$, which is proportional to the square modulus of the force and hence of a manageable complexity.

This Suzuki-Chin (SC) factorization results in a Hamiltonian

$$H^{(SC)} = \sum_{j=1}^P \frac{p_j^2}{2m} + \frac{m}{2} \left(\frac{\beta \hbar}{P} \right)^2 (q_j - q_{j+1})^2 + \tilde{V}_j(q_j, \beta) \quad (3.2)$$

which contains a bead-dependent modified potential

$$\tilde{V}_j(q_j, \beta) = w_j V(q_j) + \frac{w_j d_j}{m} \left(\frac{\beta \hbar}{P} \right)^2 |V'(q_j)|^2. \quad (3.3)$$

The coefficients w_j and d_j take different values for odd and even beads, namely $w_{2j} = 2/3$, $w_{2j+1} = 4/3$, $d_{2j} = \alpha/6$, $d_{2j+1} = (1 - \alpha)/12$, where α is a free parameter which can be varied in the range $[0, 1]$. While tuning α can improve the convergence in specific cases [16, 17], there is no universally optimal choice, so in the present work we used $\alpha = 0$.

The modified Hamiltonian resulting from the SC factorization contains the square modulus of the physical force. A path integral molecular dynamics (PIMD) calculation based on the corresponding classical equations of motion therefore requires one to evaluate the derivative of the squared force term, which contains the Hessian. Unfortunately, this is prohibitively expensive to evaluate for large systems with all but the simplest inter-atomic potentials. One way around this difficulty is to exploit the possibility of sampling based on a quasi-Trotter Hamiltonian which resembles Eq. (3.1) but has the physical potential on each bead weighted by w_j [16]. The statistics corresponding to the SC Hamiltonian can then be recovered by re-weighting configurations with the weighting factor $w = e^{-h} \equiv \exp(-\beta \Delta H/P)$, where the difference Hamiltonian has the form

$$\Delta H = \sum_{j=1}^P \frac{w_j d_j}{m} \left(\frac{\beta \hbar}{P} \right)^2 |V'(q_j)|^2. \quad (3.4)$$

The effect that re-weighting configurations based on (3.4) has on the statistical efficiency of sampling is the fundamental drawback of this approach, as will be discussed in the next subsection, based on the general results of Section 2.

(b) *The curse of system size*

The high-order path integral strategy we have just described may appear to be superior to Trotter PI in all respects [16]. However, one must not overlook the question of statistical efficiency, which could imply that longer trajectories are required to obtain the same level of sampling accuracy, a factor which must be considered when comparing the computational viability of different approaches.

Given the analysis in Section 2, all we need to evaluate the errors due to finite sampling in this context is an estimate of the variance of ΔH . Consider first the case of a simple harmonic oscillator with frequency ω , in the strongly-quantum regime, where the distribution of configurations is well approximated by the ground-state density $P(q) \propto \exp(-m\omega q^2/\hbar)$. Then, the squared fluctuations of $|V'(q)|^2$ read

$$\langle |V'(q)|^4 \rangle - \langle |V'(q)|^2 \rangle^2 = \hbar^2 m^2 \omega^6 / 2.$$

Under the simplifying assumption that the statistics of q for a finite- P path integral simulation are the same as those of the quantum oscillator at full convergence, one can write down the squared fluctuations of the temperature-scaled difference Hamiltonian $h \equiv \beta \Delta H/P$ as

$$\sigma^2(\beta \Delta H/P) \approx (P/2)^k (\beta \hbar \omega / P)^6 / 162, \quad (3.5)$$

where k varies between one (if one assumes that there is no correlation between the forces on different beads) and two (if one assumes that the forces on all beads are the same).

Eq. (3.5) implies that it is possible to reduce the fluctuations of the weights $w = e^{-h}$ at will by increasing the number of replicas. However, one should remember that keeping this number low was the primary reason for attempting a high-order factorization in the first place. Since the relative error in the energy due to the discretization of the fourth order PI grows as $(\beta\hbar\omega/P)^4$, it is clear that – for a given discretization error – the sampling problems inherent in re-weighted SC PIMD are bound to get worse as the quantum nature of the problem becomes more pronounced.

Moreover, the issue of statistical efficiency is exacerbated once one considers a system with f degrees of freedom. Since the difference Hamiltonian is size-extensive, its variance will scale linearly with f , leading to an exponential dependence of the statistical and systematic finite-sampling errors on the size of the system. For a multi-dimensional harmonic system with frequencies ω_i one can estimate – in the worst-case scenario of complete correlation of the forces on different replicas – that the total squared fluctuations will be of the order of

$$\sigma^2(\beta\Delta H/P) \approx \left(\frac{\beta\hbar\bar{\omega}}{P}\right)^6 \frac{fP^2}{648}, \quad \bar{\omega} = \left(\frac{1}{f} \sum_{i=1}^f \omega_i^6\right)^{1/6}. \quad (3.6)$$

In practice, this means that the high-order path integral integration scheme which we have described in the previous section will yield a small statistical uncertainty up to a critical system size, above which the statistical efficiency will decrease exponentially. Our analysis of the harmonic limit yields an estimate of this critical size for a given value of $\beta\hbar\bar{\omega}$. For instance, flexible water can be simulated successfully by Trotter PI with $P \gtrsim 32$, and has $\bar{\omega} \approx 2700 \text{ cm}^{-1}$. For $T = 300 \text{ K}$ and $P = 32$, Eq. (3.6) suggests that the fluctuations in the difference Hamiltonian will be of the order of one in a simulation of a few tens of water molecules. This is the sort of system size up to which re-weighted SC PIMD is therefore predicted to be advantageous.

(c) *An example: liquid water*

In order to verify the predictions of Eq. (3.6) and those of Section 2 for a realistic condensed-phase example, we have examined the case of a flexible water model [23], simulated at a temperature of 298 K and at the experimental density. We chose to perform simulations with 216 water molecules, so as to be in a regime in which – according to the arguments of Section 3b – the statistical inefficiency of re-weighting should be apparent.

We explored numbers of beads ranging between $P = 4$ and $P = 128$, and for each set of parameters performed 16 independent simulations, each of which comprised 20 ps for the initial equilibration and 1 ns during which averages were accumulated. The calculations were accelerated using the ring-polymer contraction technique [24], with the same parameters as we have employed in previous simulations [22]. Canonical sampling was enforced by applying targeted Langevin thermostats to the internal modes of the ring polymer, and a global stochastic velocity rescaling with a time constant of 10 fs to the centroid [25, 26]. The quantum mechanical kinetic energy of the water molecules was accumulated using the modified form of the centroid virial estimator proposed by Yamamoto [17]. This estimator, which is based on a scaling of fluctuation coordinates [27], can be evaluated by finite differences thereby avoiding the need for high-order derivatives of the potential.

As an initial test, we examined the convergence of the average kinetic and potential energies of the liquid as a function of the number of replicas, for both Trotter PI and




Figure 2. The two panels show the convergence of the average potential and kinetic energy for a simulation of 216 water molecules as a function of the number of beads P . The reference value – shown as a black dashed line – corresponds to Trotter PI with $P = 128$. The results from the PI-GLE accelerated Trotter PIMD are also reported [22]. Error bars were calculated, but are smaller than the data points on this scale. For the SC simulations we also report the averages of the observables computed without weights (joined by dashed lines) and those obtained from the infinite n limit of Eq. (2.4) (continuous lines).

re-weighted SC PI (see Figure 2). As expected, and as witnessed in previous calculations with a similar water model [16], the higher order of the PI factorization does translate into faster convergence. However, for the simulations with $P \lesssim 32$, the statistical error in the mean for the re-weighted SC method was found to be considerably larger than that of Trotter PIMD. Even more worrying, for re-weighted SC PIMD there is a significant difference between the weighted average of the kinetic energy and the asymptotic result predicted for Gaussian statistics [the $n \rightarrow \infty$ limit of Eq. (2.4)]. These issues will be considered in more detail below. As a final remark, it can be seen from Figure 2 that the recently-developed stochastic PI+GLE [22] approach manages to obtain even faster convergence than SC PIMD, without any of the sampling issues that we are about to elucidate.

It remains to be verified whether the scenario presented in Section 3b based on a harmonic model and an analysis of the sampling efficiency in the Gaussian limit corresponds to the numerical finding in this real example. First of all, one needs to verify whether the scaling of the variance of the logarithm of the weights follows Eq. (3.6). Figure 3 demonstrates that both the scaling with P and with the number of degrees of freedom correspond to the predictions we have made, even if the quantitative values for $\sigma^2(\beta\Delta H/P)$ are lower by about a factor of 3 than the analytical estimates obtained from the harmonic model. Taking this correction into account increases to about 100

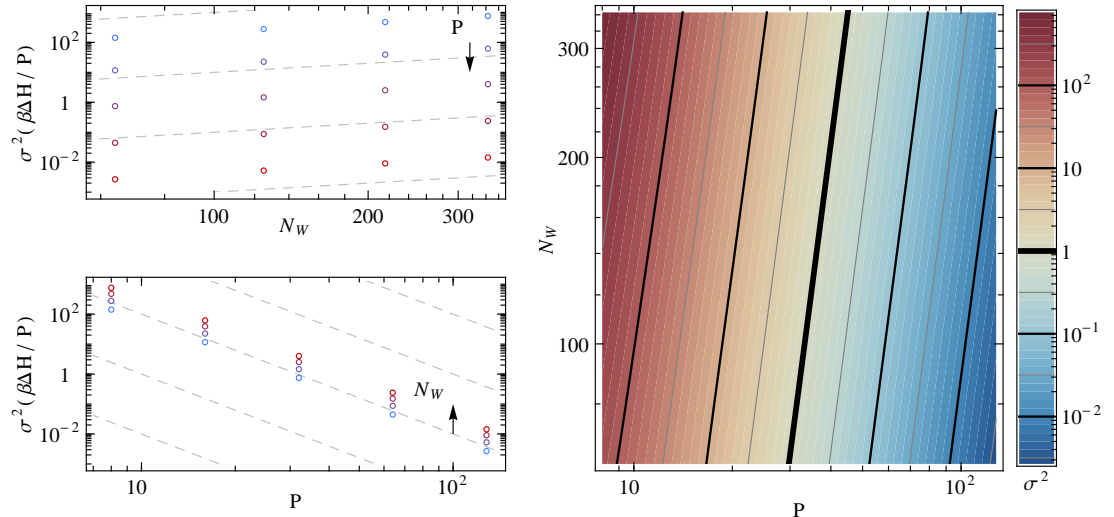


Figure 3. The three panels show the fluctuations of the temperature-scaled difference Hamiltonian $h = \beta\Delta H/P$ for a simulation of a flexible water model at $T = 298$ K, for different numbers of replicas P and numbers of water molecules N_W . The top left panel shows the data points as a function of N_W ; the dashed lines are guides for the eye corresponding to $\sigma^2 \propto N_W$ behaviour. The bottom left panel shows the points as a function of P ; here the dashed lines correspond to $\sigma^2 \propto P^{-4}$. The right panel shows a contour plot of the interpolant of the data points as a function of P and N_W , with σ^2 denoted by the colour scale. Red indicates regions in which the statistical inefficiency of re-weighted Suzuki-Chin PIMD is unmanageable, while blue indicates regions in which it is comparable to that of Trotter PIMD. Note that for a given P , the inefficiency gets worse with increasing system size.

the number of water molecules for which re-weighted SC PIMD is computationally reasonable.

Having verified that the fluctuations in the difference Hamiltonian closely obey Eq. (3.6), let us discuss the validity of the relations between these fluctuations and the error in the mean that we have worked out in Section 2. Figure 4 shows that – for the case of $P = 16$, where the variance of the temperature-scaled difference Hamiltonian $h = \beta\Delta H/P$ is much larger than one – the trends do indeed correspond to those derived in the limit of Gaussian statistics. Averages computed from short trajectories suffer a large systematic bias, and they slowly approach the limit predicted from the cross-correlation between the observable and the difference Hamiltonian as the number of samples is increased. To obtain a converged expectation value and therefore verify the accuracy of the prediction based on Gaussian statistics, a prohibitively long trajectory would be required. Also as anticipated, the upper panel of Figure 4 clearly demonstrates that the statistical variance in the mean for re-weighted SC PIMD decreases much more slowly than for Trotter PIMD, completely offsetting the benefits of using the higher-order propagator.

It is important to remark that we have been able to make these issues stand out clearly because we have considered a relatively simple problem, for which we could run long trajectories and gather extensive statistics. In real applications one can rarely afford to perform such a detailed error analysis, and in this respect the analytical results we have derived provide a useful litmus test. The fluctuations of the temperature-scaled




Figure 4. The two panels show the convergence of the mean kinetic energy for a simulation of 216 water molecules, at $T = 298$ K and for $P = 16$, as a function of the length of the simulation, together with the relative fluctuations in the mean. In the lower plot, the un-weighted average kinetic energy and that obtained from the infinite n limit of Eq. (2.4) are also shown.

difference Hamiltonian h should be smaller than one, since for larger fluctuations the statistical inefficiency rapidly becomes unmanageable. Even more interesting is the cross-correlation between the observable a and h , which indicates the extent of the systematic bias.

4. Conclusions

In this paper we have performed a careful assessment of the statistical efficiency of re-weighted sampling. Our analysis shows a dramatic degradation in performance when the squared fluctuations of the logarithm of the weight exceed one. While it is common knowledge that large fluctuations in the weights lead to failure of such techniques, and that the size-extensivity of the difference Hamiltonian makes these methods unsuitable for the study of large-scale problems, our analysis shows clearly how dramatic this effect can be, and also that the problem is quite insidious.

In fact, insufficient sampling in a re-weighted calculation is very difficult to detect. It mainly manifests itself through the error in the mean decreasing more slowly than the inverse square root of the number of samples even after a very large number of samples have been accumulated, indicating that the asymptotic regime has not yet been reached (see the upper panel of Fig. 4). An even greater concern is that finite sampling implies a bias in the estimate of the observable, which decreases very slowly and would be

exceedingly hard to detect in a more complex – or computationally demanding – problem (see the lower panel of Fig. 4).

As a demonstration of how these results can be used, we have examined the applicability of high order path integral factorizations to condensed-phase simulations. The most promising of these techniques relies on re-weighting based on a difference Hamiltonian [16, 17]. We have worked out the dependence of the fluctuations in this quantity on the properties of the system being studied, the number of degrees of freedom involved, and the number of path integral beads. We have shown that for large and/or strongly quantum systems, a large number of replicas is required, not so much in order to converge the asymptotic expectation values of quantum observables, but to keep the fluctuations in the re-weighted average under control. Our analytical results allow one to predict the break-even point at which the advantages of high-order path integrals are lost due to statistical inefficiency. These predictions have been qualitatively verified by numerical experiments on the simulation of a flexible model of water under ambient conditions. For this system, high-order SC PIMD is advantageous for simulations of up to around 100 water molecules, but no more.

While a niche for the use of Hessian-free high-order PIMD therefore exists for the study of small clusters, one must consider that accelerated PIMD based on stochastic dynamics [22] exhibits comparable or faster convergence of quantum observables, without being affected by any of the statistical issues we have focussed on here. We firmly believe that this stochastic acceleration is therefore a far more promising approach to large scale path integral simulations.

Acknowledgment

The authors gratefully acknowledge an insightful discussion with Michele Parrinello, and funding from the Wolfson Foundation, the Royal Society, and the Swiss National Science Foundation.

An asymptotic expression for re-weighted averages

In this appendix we present a derivation of the asymptotic expansion of the expectation value of a re-weighted average of n terms from a general distribution. As usual, we say that the formal sum $\sum_{j=0}^{\infty} c_j n^{-j}$ is an asymptotic expansion for f_n if for any fixed k we have $f_n = \sum_{j=0}^k c_j n^{-j} + \mathcal{O}(n^{-(k+1)})$ as $n \rightarrow \infty$. This does not imply that the sum converges. Our main result is as follows:

THEOREM 1. *Let $(a_i, w_i)_{i=1}^n$ be n independent samples from a (correlated) joint distribution $p_0(a, w)$. Suppose that w takes only positive values, and that all moments of a , w and w^{-1} are finite. Let*

$$\bar{a}_n = \frac{\sum_{i=1}^n a_i w_i}{\sum_{i=1}^n w_i}. \quad (\text{A.1})$$

Then $\langle \bar{a}_n \rangle$ and $\sigma^2(\bar{a}_n) = \langle \bar{a}_n^2 \rangle - \langle \bar{a}_n \rangle^2$ have asymptotic expansions in n whose coefficients are determined by the joint moments of $p_0(a, w)$, with the leading terms in these expansions being given by Eqs. (2.2) and (2.3).

The argument below will show how to calculate the higher-order coefficients of the expansion, although there does not seem to be a simple explicit formula for the general coefficient c_j .

Multiplying the distribution w by a constant (i.e., multiplying all w_i by the same constant) does not affect $\langle \bar{a}_n \rangle$, so we assume in the proof that $\langle w \rangle = 1$. With this assumption, Eqs. (2.2) and (2.3) simplify to

$$\langle \bar{a}_n \rangle = \langle aw \rangle + n^{-1} (\langle aw \rangle \langle w^2 \rangle - \langle aw^2 \rangle) + \mathcal{O}(n^{-2}), \quad (\text{A.2})$$

$$\sigma^2(\bar{a}_n) = n^{-1} \left(\langle a^2 w^2 \rangle - 2 \langle aw \rangle \langle aw^2 \rangle + \langle aw \rangle^2 \langle w^2 \rangle \right) + \mathcal{O}(n^{-2}). \quad (\text{A.3})$$

Explicit calculations can be further simplified by subtracting a constant from a (and hence from $\langle \bar{a}_n \rangle$) to ensure that $\langle aw \rangle = 0$, but we shall not do this in the following.

Since all the samples are independent and identically distributed, one can write

$$\langle \bar{a}_n \rangle = n \left\langle \frac{a_n w_n}{z_n} \right\rangle = n \left\langle \frac{aw}{w + z_{n-1}} \right\rangle \quad (\text{A.4})$$

where $z_n = \sum_{i=1}^n w_i$ is the sum of n independent copies of w , and z_{n-1} (but not z_n) is independent of the correlated pair (a, w) .

Let $\tilde{w} = w - 1$ and $\tilde{z}_n = z_n - n = \sum_{i=1}^n \tilde{w}_i$ be the centralized versions of w and z_n . We first consider the properties of \tilde{z}_n . Even in the (canonical) special case where w is log-normal, not much can be obtained analytically about the distribution of \tilde{z}_n [28]. However, one can obtain its moments using the general formula

$$\kappa^{(k)} = \mu^{(k)} - \sum_{i=1}^{k-1} \binom{k-1}{i-1} \kappa^{(i)} \mu^{(k-i)} \quad (\text{A.5})$$

relating the cumulants $\kappa^{(k)}$ and moments $\mu^{(k)}$ of a distribution. When $\mu^{(1)} = 0$ (and hence $\kappa^{(1)} = 0$), as is the case for the centred variables \tilde{w} and \tilde{z}_n , Eq. (A.5) can be written as

$$\tilde{\kappa}^{(k)} = \tilde{\mu}^{(k)} - \sum_{i=2}^{k-2} \binom{k-1}{i-1} \tilde{\kappa}^{(i)} \tilde{\mu}^{(k-i)}. \quad (\text{A.6})$$

We therefore first calculate the cumulants $\tilde{\kappa}_1^{(k)}$ of the distribution of \tilde{w} using Eq. (A.6). The cumulants of \tilde{z}_n are simply $\tilde{\kappa}_n^{(k)} = n \tilde{\kappa}_1^{(k)}$, and the moments $\tilde{\mu}_n^{(k)}$ of \tilde{z}_n can be recovered by inverting Eq. (A.6). By induction on k we find that $\tilde{\mu}_n^{(k)}$ is a polynomial in n of degree $\lfloor k/2 \rfloor$; this can also be seen directly by expanding the expectation of the k th power of the sum of the \tilde{w}_i .

An asymptotic expansion for $\langle \bar{a}_n \rangle$ can now be developed by expanding the denominator in Eq. (A.4) about its mean. Let $X = \tilde{z}_n/n = z_n/n - 1$. Using the identity $1/(1+x) = \sum_{j=0}^{k-1} (-x)^j + (-x)^k/(1+x)$, we have

$$\begin{aligned} \langle \bar{a}_n \rangle &= \left\langle \frac{a_n w_n}{z_n/n} \right\rangle = \left\langle \frac{a_n w_n}{1+X} \right\rangle = \\ &= \sum_{j=0}^{k-1} (-1)^j \langle a_n w_n X^j \rangle + (-1)^k \left\langle a_n w_n \frac{X^k}{1+X} \right\rangle. \end{aligned} \quad (\text{A.7})$$

Since $X = (\tilde{z}_{n-1} + \tilde{w}_n)/n$ and \tilde{z}_{n-1} is independent of (a_n, w_n) , the terms in the sum can be written out in terms of the moments $\tilde{\mu}_{n-1}^{(k)}$ of \tilde{z}_{n-1} :

$$(-1)^j \langle a_n w_n X^j \rangle = (-n)^{-j} \sum_{i=0}^j \binom{j}{i} \langle a(\tilde{w} + 1)\tilde{w}^i \rangle \tilde{\mu}_{n-1}^{(j-i)}. \quad (\text{A.8})$$

For each j , the term described by Eq. (A.8) is a (Laurent) polynomial in n with powers between $n^{-\lceil j/2 \rceil}$ and n^{-j} . We claim that

$$R_n^{(k)} = \left\langle a_n w_n \frac{X^k}{1+X} \right\rangle = \mathcal{O}(n^{-k/2}), \quad (\text{A.9})$$

so the residual term $R_n^{(k)}$ shrinks with n . Moreover, summing Eq. (A.8) over $j \leq 2d$ and collecting coefficients gives the first terms of the asymptotic expansion of $\langle \bar{a}_n \rangle$ down to order n^{-d} . For example, setting $d = 1$ and discarding the terms with an n^{-2} dependence yields Eq. (A.2).

The bound Eq. (A.9) is perhaps most easily seen by applying the Cauchy–Schwartz inequality twice, or Hölder’s inequality once, to show that in general $|\langle A_1 A_2 A_3 A_4 \rangle| \leq \prod_{i=1}^4 \langle A_i^4 \rangle^{1/4}$. Here $A_1 = a_n$ and $A_2 = w_n$ contribute constant factors to the product. With $A_3 = X^k = (\tilde{z}_n/n)^k$ we have $\langle A_3^4 \rangle = n^{-4k} \tilde{\mu}_n^{(4k)} = \mathcal{O}(n^{-2k})$, so $\langle A_3^4 \rangle^{1/4}$ has the required order $\mathcal{O}(n^{-k/2})$. Finally, $A_4 = (1+X)^{-4} = (z_n/n)^{-4}$. Since the function x^{-4} is convex, and z_n/n is just the average of the w_i , we have $A_4^4 \leq \sum_{i=1}^n w_i^{-4}/n$, so $\langle A_4^4 \rangle \leq \langle w^{-4} \rangle$, which is by assumption a finite constant, and the bound Eq. (A.9) follows.

For the variance, the argument is very similar, now using the fact that

$$\langle \bar{a}_n^2 \rangle = n(n-1) \left\langle \frac{a_{n-1} w_{n-1} a_n w_n}{n^2 (1+X)^2} \right\rangle + n \left\langle \frac{a_n^2 w_n^2}{n^2 (1+X)^2} \right\rangle,$$

and the identity $(1+x)^{-2} = \sum_{j=0}^{k-1} (j+1)(-x)^j + (-1)^k \frac{(k+1)x^k + kx^{k+1}}{(1+x)^2}$. A similar method applies to higher moments of \bar{a}_n .

Note that the assumption that all negative moments of w are finite was only used to control the tail behaviour of $1/(1+X) = (z_n/n)^{-1}$ near 0. In fact, this has much better tail behaviour than w^{-1} , so much weaker conditions suffice. For example, splitting the sum z_n into groups of m and noting that $z_m^{-1} \leq \prod_{i=1}^m w_i^{-1/m}$, one can check that the result holds assuming only that $\langle w^{-\alpha} \rangle$ is finite for *some* positive α .

Finally, we should emphasize that one cannot expect the infinite sum corresponding to Eq. (A.7), or the asymptotic expansion of $\langle \bar{a}_n \rangle$, to converge for any fixed n . Indeed, in the log-normal case where $w = e^{-h}$, we have $\mu_1^{(k)} = \langle w^k \rangle = \exp(\sigma^2(h)k^2/2)$, which becomes $\exp(\sigma^2(h)k(k-1)/2)$ after normalizing so that $\langle w \rangle = 1$. Moreover, the k th central moment $\tilde{\mu}_1^{(k)}$ also contains a component proportional to $\mu_1^{(k)}$ (it is a linear combination of this and lower moments). Hence, Eq. (A.8) will contain a component proportional to $[\exp(\sigma^2(h)/2)]^{j(j-1)}/n^{j-1}$. For any n , this eventually increases with j , so the infinite sum does not converge. Similar considerations strongly suggest that the first few terms in the asymptotic expansion will be a good approximation only if $\sigma^2(h)$ is small compared to (at most a certain constant times) $\log n$.

References

- [1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.*, 21(6):1087–1092, 1953.
- [2] W. K. Hastings. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*, 57(1):–97, 1970.
- [3] D. Frenkel and B. Smit. *Understanding Molecular Simulation*. Academic Press, London, Second edition, 2002.
- [4] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Phys. Lett. B*, 195(2):216–222, 1987.
- [5] S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman. The weighted histogram analysis method for free-energy calculations on biomolecules. I. The method. *J. Comp. Chem.*, 13(8):1011–1021, 1992.
- [6] M. Bonomi, A. Barducci, and M. Parrinello. Reconstructing the equilibrium Boltzmann distribution from well-tempered metadynamics. *J. Comp. Chem.*, 30(11):1615–1621, 2009.
- [7] I. M. Barbour, S. E. Morrison, E. G. Klepfish, J. B. Kogut, and M. P. Lombardo. Results on finite density QCD. *Nucl. Phys. B*, 60(1-2):220–233, 1998.
- [8] Z. Fodor and S. D. Katz. A New method to study lattice QCD at finite temperature and chemical potential. *Phys. Lett. B*, 534(1-4):87–92, 2002.
- [9] A. F. M. Smith and G. O. Roberts. Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods. *J. Royal Stat. Soc. B*, 55(1):3–23, 1993.
- [10] R. P. Feynman and A. R. Hibbs. *Quantum Mechanics and Path Integrals*. McGraw-Hill, New York, 1964.
- [11] M. Parrinello and A. Rahman. Study of an F center in molten KCl. *J. Chem. Phys.*, 80:860, 1984.
- [12] D. M. Ceperley. Path integrals in the theory of condensed helium. *Rev. Mod. Phys.*, 67(2):279–355, Apr 1995.
- [13] M. Takahashi and M. Imada. Monte Carlo calculation of quantum systems. II. Higher order correction. *J. Phys. Soc. Jap.*, 53(11):3765–3769, 1984.
- [14] M. Suzuki. Hybrid exponential product formulas for unbounded operators with possible applications to Monte Carlo simulations. *Phys. Lett. A*, 201(5-6):425–428, 1995.
- [15] S. A. Chin. Symplectic integrators from composite operator factorizations. *Phys. Lett. A*, 226(6):344–348, 1997.
- [16] S. Jang, S. Jang, and G. A. Voth. Applications of higher order composite factorization schemes in imaginary time path integral simulations. *J. Chem. Phys.*, 115(17):7832–7842, 2001.

- [17] T. M. Yamamoto. Path-integral virial estimator based on the scaling of fluctuation coordinates: Application to quantum clusters with fourth-order propagators. *J. Chem. Phys.*, 123:104101, 2005.
- [18] L. S. Schulman. *Techniques and applications of path integration*. John Wiley, 1996.
- [19] D. Chandler and P. G. Wolynes. Exploiting the isomorphism between quantum theory and classical statistical mechanics of polyatomic fluids. *J. Chem. Phys.*, 74(7):4078–4095, 1981.
- [20] S. A. Chin. A fundamental theorem on the structure of symplectic integrators. *Phys. Lett. A*, 354(5-6):373–376, 2006.
- [21] X. P. Li and J. Q. Broughton. High-order correction to the Trotter expansion for use in computer simulation. *J. Chem. Phys.*, 86:5094, 1987.
- [22] M. Ceriotti, D. E. Manolopoulos, and M. Parrinello. Accelerating the convergence of path integral dynamics with a generalized Langevin equation. *J. Chem. Phys.*, 134(8):084104, 2011.
- [23] S. Habershon, T. E. Markland, and D. E. Manolopoulos. Competing quantum effects in the dynamics of a flexible water model. *J. Chem. Phys.*, 131(2):024501, 2009.
- [24] T. E. Markland and D. E. Manolopoulos. A refined ring polymer contraction scheme for systems with electrostatic interactions. *Chem. Phys. Lett.*, 464(4-6):256, 2008.
- [25] M. Ceriotti, M. Parrinello, T. E. Markland, and D. E. Manolopoulos. Efficient stochastic thermostating of path integral molecular dynamics. *J. Chem. Phys.*, 133(12):124104, 2010.
- [26] G. Bussi, D. Donadio, and M. Parrinello. Canonical sampling through velocity rescaling. *J. Chem. Phys.*, 126(1):014101, 2007.
- [27] C. Predescu, D. Sabo, J. D. Doll, and D. L. Freeman. Heat capacity estimators for random series path-integral methods by finite difference schemes. *J. Chem. Phys.*, 119(23):12119–12128, 2003.
- [28] L. Fenton. The sum of log-normal probability distributions in scatter transmission systems. *IRE Trans. Comm. Sys.*, 8(1):57–67, 1960.