

Improving dynamical lattice QCD simulations through integrator tuning using Poisson brackets and a force-gradient integrator

M. A. Clark,¹ Bálint Joó,² A. D. Kennedy,³ and P. J. Silva^{4,*}

¹*Harvard-Smithsonian Center for Astrophysics, Cambridge, MA 02138, U.S.A.*

²*Jefferson Lab, 12000 Jefferson Avenue, Newport News, VA 23606, U.S.A.*

³*Tait Institute and SUPA, School of Physics & Astronomy,
The University of Edinburgh, Edinburgh EH9 3JZ, Scotland, U.K.*

⁴*Centro de Física Computacional, Universidade de Coimbra, Portugal*

(Dated: August 8, 2011)

We show how the integrators used for the molecular dynamics step of the Hybrid Monte Carlo algorithm can be further improved. These integrators not only approximately conserve some Hamiltonian H but conserve exactly a nearby shadow Hamiltonian \tilde{H} . This property allows for a new tuning method of the molecular dynamics integrator and also allows for a new class of integrators (force-gradient integrators) which is expected to reduce significantly the computational cost of future large-scale gauge field ensemble generation.

INTRODUCTION AND MOTIVATION

Hybrid Monte Carlo (HMC) [1] is the algorithm of choice to generate lattice QCD configurations including the effect of dynamical fermions. The most time consuming ingredient of HMC is the molecular dynamics (MD) step, which consists of a reversible volume-preserving approximate MD trajectory of $\tau/\delta\tau$ steps (with τ being the length of the trajectory and $\delta\tau$ the stepsize) followed by a Metropolis accept/reject test with acceptance probability $\min(1, e^{-\delta H})$ where δH is the change in the Hamiltonian $H = T + S$ whose kinetic and potential parts are T and S .

A molecular dynamics trajectory is not only an approximate integral curve of the Hamiltonian vector field \hat{H} corresponding to H , but is also an exact integral curve of the Hamiltonian vector field $\hat{\tilde{H}}$ of an exactly conserved shadow Hamiltonian \tilde{H} . The asymptotic expansion of this shadow Hamiltonian in the stepsize $\delta\tau$ may be computed using the Baker–Campbell–Hausdorff (BCH) formula and expressed in terms of Poisson brackets (PBs) [2, 3]. As a simple example consider the PQQQP (also known as 2MN [4]) integrator

$$U_{\text{PQQQP}}(\tau) = \left(e^{\lambda \hat{S} \delta\tau} e^{\frac{1}{2} \hat{T} \delta\tau} e^{(1-2\lambda) \hat{S} \delta\tau} e^{\frac{1}{2} \hat{T} \delta\tau} e^{\lambda \hat{S} \delta\tau} \right)^{\tau/\delta\tau}$$

whose shadow Hamiltonian is

$$\begin{aligned} \tilde{H}_{\text{PQQQP}} = & H + \left(\frac{6\lambda^2 - 6\lambda + 1}{12} \{S, \{S, T\}\} + \frac{1 - 6\lambda}{24} \{T, \{S, T\}\} \right) \delta\tau^2 \\ & + \left(\frac{-1 + 30\lambda^2 - 60\lambda^3 + 30\lambda^4}{720} \{S, \{S, \{S, \{S, T\}\}\}\} + \frac{-4 + 15\lambda + 15\lambda^2 - 30\lambda^3}{720} \{T, \{S, \{S, \{S, T\}\}\}\} \right. \\ & + \frac{-7 + 30\lambda}{1440} \{T, \{T, \{S, \{S, T\}\}\}\} + \frac{-7 + 30\lambda}{5760} \{T, \{T, \{T, \{S, T\}\}\}\} \\ & \left. + \frac{-2 + 15\lambda - 35\lambda^2 + 30\lambda^3}{240} \{\{S, T\}, \{S, \{S, T\}\}\} + \frac{-2 + 15\lambda - 30\lambda^2}{720} \{\{S, T\}, \{T, \{S, T\}\}\} \right) \delta\tau^4 + \mathcal{O}(\delta\tau^6). \end{aligned} \quad (1)$$

Note that we have one free parameter, λ , which is often set to some value not taking PBs into account. In [5], the authors chose λ by minimising δH empirically, requiring a sequence of runs at different values of λ . Others used $\lambda_c \approx 0.193183$ [4], which minimizes the norm of the coefficients of the PBs in the second-order term. However, this is not necessarily the best choice.

We have evaluated PBs and shadow Hamiltonians for gauge theories (where gauge fields are constrained to live on a Lie group manifold) for the first time [3, 6–8]. There-

fore, in this letter we propose to measure the volume-averaged PBs and tune the free parameters of an MD integrator taking PB measurements into account. As we will see, our tuning procedure also allows us to find out the best number of steps of a nested integrator scheme. We also present a new integrator step and a new integrator which will be able to reduce the cost of large volume simulations.




FIG. 1. Comparison of measured acceptance rates and their predictions from Poisson bracket measurements. In the left hand plot we fix $\delta\tau = 0.1$ and leave λ as a free parameter, whereas in the right we take $\lambda = 0.18$ and plot P_{acc} as a function of $\delta\tau$.

INTEGRATOR TUNING

Let us define the difference between the shadow (\tilde{H}) and actual (H) Hamiltonians as $\Delta H = \tilde{H} - H$. Noting that $\text{Var}(\Delta H)$ means the variance of the distribution of values of ΔH over phase space, one can show that the acceptance rate P_{acc} can be given by [8]

$$P_{\text{acc}} = \text{erfc} \left(\sqrt{\frac{1}{4} \text{Var}(\Delta H)} \right). \quad (2)$$

To estimate P_{acc} from eq. 2, one only needs to measure the PB from equilibrated configurations. This allow us to express P_{acc} as a function of the integrator parameters and find their optimal values that maximize P_{acc} .

As a simple test, we consider an HMC simulation of two flavors of Wilson fermions at $\kappa = 0.158$ and Wilson gauge action at $\beta = 5.6$ on an 8^4 lattice. We use a single level PQQQP integrator and a unit trajectory length, therefore we have two tunable parameters: the integrator parameter λ and the step size $\delta\tau$. We measure ΔH up to fourth order in $\delta\tau$. In Figure 1 we compare the acceptance rates predicted by the formula above with numerical data taken from simulations at various values of λ and $\delta\tau$. The PB values used for the predictions were measured at $\lambda = 0.18$ and $\delta\tau = 0.1$ – but one should note that our predictions are independent of the integrator parameters used to get the Poisson bracket values.

The plots show good agreement between predicted and measured acceptance rates, provided the stepsize is not too big, otherwise the BCH expansion breaks down. Moreover, the maximum of the acceptance rate in the left hand plot is achieved at $\lambda_{\text{max}} \approx 0.1836$ (to be compared with λ_c). We now use eq. (2) to tune the MD integrator on a larger volume. Ultimately, we are interested in reducing the computational cost, which depends

on the wall-clock time spent computing the force terms on a unit of MD time as well as the acceptance rate, and the autocorrelation time τ_{corr} for the observables. We neglect τ_{corr} in this discussion as it is not sensitive to the choice of integrator parameters as long as the acceptance rate is reasonable, and define our cost metric as

$$\text{cost} = \frac{\text{trajectory CPU time}}{P_{\text{acc}} \tau}. \quad (3)$$

For our purposes, the numerator of eq. (3) is estimated by considering the time spent in force computation along the trajectory. In particular, for a nested integrator, the numerator of this cost function is a function of the number of steps at each level times the CPU time required to compute the forces at that level. Therefore, minimizing eq. (3) will allow us to both find out the optimal integrator parameters, as well as find out the optimal stepsize, or the number of steps at each level of a nested integrator scheme. This is a more direct approach than the popular “balancing forces” method [9].

TUNING A REAL SIMULATION

Level i	Force	F time	FG time
0	Hasenbusch ($\mu = 0 / \mu = 0.057$)	21.21 s	26.61 s
1	Hasenbusch ($\mu = 0.057 / \mu = 0.25$)	3.98 s	7.55 s
2	Wilson ($\mu = 0.25$)	1.05 s	1.98 s
3	Gauge	0.075 s	0.142 s

TABLE I. Set-up used in the HMC simulation described in this section, together with typical times spent on force computation. μ is the twisted mass parameter [9]. For convenience, times for the force-gradient computation are also shown here.

As an application of our tuning technology, we consider a HMC simulation of a $24^3 \times 32$ lattice, with two flavours of Wilson fermions with $\kappa = 0.1580$ and $\beta = 5.6$. As in [9], we use a nested PQQQP integrator scheme, with the inclusion of two Hasenbusch fields with twisted mass fermions as “preconditioners”. In [9], each nested level of the integrator has one force term, and the free parameter of the PQQQP integrator has been set to $\lambda = 1/6$ at all levels. In Table I, for each level i (note that 0 is the outermost level), we show the type of force and its parameters, and mean values of the time spent on force and force-gradient computation.

In order to improve the integrator scheme used in [9], we considered two different nested schemes:

- PQ4. the original scheme, but with tuned values of λ ;
- PQ3. the two Hasenbusch fields appear now at the same level (so we have only 3 different levels).

Tuning Scheme	m_i				λ_i				Prediction		Measurement		
	0	1	2	3	0	1	2	3	P_{acc}	F time / traj.	P_{acc}	Time / traj.	Cost
Original	3	1	2	3	1/6	1/6	1/6	1/6	0.85(1)(3)	308 s	0.88(4)	405 s	463
PQ4 / 4th	3	1	2	1	0.1903(36)	0.1696(66)	0.1885(69)	0.1670(80)	0.89(1)(0)	294 s	0.85(3)	399 s	471
PQ4 / 4th	3	1	1	2	0.1966(64)	0.1660(131)	0.1885(35)	0.1524(168)	0.80(2)(0)	267 s	0.82(5)	360 s	438
PQ3 / 4th	3	3	2	—	0.1803(22)	0.1902(53)	0.1281(220)	—	0.83(3)(0)	234 s	0.83(4)	345 s	417
PQ3 / 2nd	3	3	2	—	0.1735(25)	0.1924(53)	0.1415(216)	—	0.81(2)(3)	234 s	0.72(4)	354 s	491
PQ3	3	3	2	—	λ_c	λ_c	λ_c	—	0.76(3)(3)	234 s	0.80(3)	339 s	425
PQ3	3	3	2	—	1/6	1/6	1/6	—	0.73(3)(5)	234 s	0.74(5)	349 s	473
PQ3 / 4th	3	3	1	—	0.1793(24)	0.1890(65)	0.1636(66)	—	0.81(2)(0)	228 s	0.78(5)	342 s	442

TABLE II. Tuning of the PQPQP integrator scheme. All errors shown are statistical, with the exception of the second set of errors in the predicted acceptance rates, which are some sort of systematic error, estimated from the difference of predicting acceptance rates using a shadow Hamiltonian up to $\delta\tau^2$ or $\delta\tau^4$. All times refer to runs utilizing 128 cores of the Iridis cluster.

Table II shows parameters which minimize the cost metric (3). In our simulations, we have fixed $\tau = 1$, unless stated otherwise. For each scheme (which is also described by the highest power of $\delta\tau$ used to compute ΔH), we show the optimal number of steps at each level, the optimal λ parameters, our predictions for the acceptance rate, the estimated time spent in force computation in one trajectory, and measurements of acceptance rates and trajectory times. For comparison we also show data for the original scheme [9].

We see that all predicted acceptance rates agree, within errors, with the measured ones. Furthermore, the tuning of λ 's in the PQ4 scheme allows a reduction of the number of steps on the inner levels, so the CPU time per trajectory decreases. Moreover, the PQ3 scheme allows further improvement in cost measures. For this scheme, we also show the performance obtained using other λ values. We conclude that, for an optimal choice of integrator parameters, one is encouraged to tune the integrator using the best available approximation to ΔH .

FORCE-GRADIENT INTEGRATOR

Since the Poisson bracket $\{S, \{S, T\}\}$ does not depend on momentum [3], we can evaluate the integrator step $e^{\{S, \widehat{\{S, T\}}\} \delta\tau^3}$ explicitly. If one uses $\lambda = 1/6$ for the PQPQP integrator together with this integrator step, we eliminate all $\mathcal{O}(\delta\tau^2)$ terms in ΔH . We therefore define a PQPQP *force-gradient* integrator as

$$U_{FG}(\tau) = \left(e^{\frac{1}{6}\hat{S}\delta\tau} e^{\frac{1}{2}\hat{T}\delta\tau} e^{\frac{48\hat{S}\delta\tau - \widehat{\{S, \{S, T\}\}} \delta\tau^3}{72}} e^{\frac{1}{2}\hat{T}\delta\tau} e^{\frac{1}{6}\hat{S}\delta\tau} \right)^{\tau/\delta\tau}.$$

Note that the performance of this integrator has been shown to be much better than Campostrini integrator [7]. This is not surprising since the coefficients of the $\mathcal{O}(\delta\tau^4)$ term in the shadow Hamiltonian are about two orders of magnitude smaller [3].

In order to test the performance of this integrator, we considered the two integrator schemes defined in the last section, replacing the PQPQP integrator, at all levels, by this new one. We will denote these new schemes as FG4 and FG3 which have 4 and 3 different levels, respectively. As in this case there are no tuneable parameters, we could only vary the number of steps at each level. In Table III we show the best parameters we found as well as the measured values of acceptance rates and trajectory times.

Scheme	τ	m_i			Prediction		Measurement			
		0	1	2	3	P_{acc}	$F + FG$ time / traj.	P_{acc}	Time / traj.	cost
FG4	1.0	3	1	1	1	0.97(1)	415 s	0.92(2)	523 s	569
FG4	1.1	3	1	1	1	0.96(1)	415 s	0.85(4)	526 s	560
FG4	1.2	3	1	1	1	0.94(1)	415 s	0.71(6)	500 s	585
FG3	1.0	3	1	1	—	0.91(1)	314 s	0.80(3)	393 s	492

TABLE III. Tuning of FG4 and FG3 integrator schemes. Note that P_{acc} predictions are not compatible with measurements. We expect this to be due to higher order corrections to ΔH .

A comparison between Tables II and III shows that the use of a force-gradient integrator allows a smaller number of steps in the inner levels. Furthermore, comparing the 4 nested level schemes, one sees that the force-gradient acceptance rates are higher. However, the trajectory CPU times also increase (because we have to evaluate one more step), so the cost measures are higher.

A LARGER VOLUME

Despite the comparison done in the last section does not favour the use of a force-gradient integrator, it is expected this integrator will be of use for larger lattice volumes. Thus we now consider a thermalized 40^4 simulation (using the same action parameters as used in the

Tuning Scheme	m_i				λ_i				Prediction		Measurement		
	0	1	2	3	0	1	2	3	P_{acc}	F time / traj.	P_{acc}	Time / traj.	Cost
	5	1	2	1	0.1940(32)	0.1712(53)	0.1961(50)	0.1757(64)	0.90(1)(0)	1819 s	0.85(4)	2143 s	2513
FG4	4	1	1	1	—	—	—	—	0.95(1)	2196 s	0.84(4)	2462 s	2931
PQ3 / 4th	5	3	1	—	0.1780(18)	0.1995(49)	0.1794(55)	—	0.83(2)(1)	1513 s	0.82(4)	1934 s	2355
FG3	4	1	1	—	—	—	—	—	0.92(1)	1816 s	0.82(5)	2158 s	2641

TABLE IV. Tuning of a 40^4 simulation. CPU times refer to runs utilizing 256 cores of the Iridis cluster.

previous sections) and proceeded with a similar tuning analysis. We show the results in Table IV. Although we were able to use a higher stepsize for the schemes which use a force-gradient integrator, the trajectory CPU times are still higher than PQQP runs.

Using the data available, one can compute ratios between the cost of the FG schemes over PQQP ones; we see they decrease for larger volumes — see Table V. From these results, we can see that at some increased volume (assuming constant physics) the force-gradient integrator will become more efficient than the PQQP integrator.

Nested scheme	FG4/PQ4	FG3/PQ3
$24^3 \times 32$	1.30	1.18
40^4	1.17	1.12

TABLE V. Cost ratios (FG over PQQP).

This cross-over point can be estimated from appealing to the requirement that the equilibrium distribution must satisfy $\langle e^{\delta H} \rangle = 1$. Expanding to second order, we have $\langle \delta H \rangle \sim \frac{1}{2} \langle \delta H^2 \rangle$, thus for second- and fourth-order integrators $\langle \delta H \rangle \sim \delta \tau^4$ and $\langle \delta H \rangle \sim \delta \tau^8$, respectively. The cost has a trivial linear volume factor and scales linearly with $1/\delta \tau$. Since δH is extensive in the volume we can equate it with the volume stepsize scaling, thus we have $\text{cost}^{\text{PQQP}} \sim V^{5/4}$ and $\text{cost}^{\text{FG}} \sim V^{9/8}$, so the cost ratio will behave as $V^{-1/8}$. Using the data in table V to estimate the multiplicative coefficient, we estimate that the force-gradient integrator becomes more efficient at $V = 56^4$ and $V = 50^4$ for the four-level and three-level integrators, respectively. Given current leading-edge lattice computations are presently at this volume, or slightly larger, we therefore predict that the force-gradient integrator will reduce significantly the computational cost for future large-scale gauge field ensemble generation.

CONCLUSIONS

We have presented a novel way of tuning a HMC integrator, together with practical examples. This tuning procedure can be used for all lattice gauge and fermionic

actions, and allows a systematic study of the MD integrators currently used in large scale dynamical lattice simulations. We have also presented a new fourth-order integrator which is expected to reduce significantly the computational cost of HMC simulations.

We are currently working towards a general implementation of the calculation of Poisson brackets and force-gradient terms in Chroma [10]. Details about computing Poisson Brackets will be presented elsewhere [11]. In the near future we will also consider tuning simulations using other lattice actions.

PJS acknowledges support from FCT via grant SFRH/BPD/40998/2007, and project PTDC/FIS/100968/2008. BJ acknowledges funding through US D.O.E Grants DE-FC02-06ER41440, DE-FC02-06ER41449 (SciDAC) and DE-AC05-06OR23177 under which Jefferson Science Associates LLC manages and operates the Jefferson Lab. MAC acknowledges support from the NSF via PHY-0835713 and OCI-1060067. The U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce this manuscript for U.S. Government purposes. The numerical results have been obtained using Chroma library [10]. Simulations have been carried out in Iridis (University of Southampton) and Centaurus (University of Coimbra) High Performance Computing facilities.

* Corresponding author: psilva@teor.fis.uc.pt

- [1] S. Duane, A. D. Kennedy, B. J. Pendleton and D. Roweth, *Phys. Lett. B* **195** (1987) 216.
- [2] M. A. Clark and A. D. Kennedy, *Phys. Rev. D* **76** (2007) 074508, [arXiv:0705.2014 \[hep-lat\]](#).
- [3] A. D. Kennedy and M. A. Clark, *PoS(LATTICE 2007)038*, [arXiv:0710.3611 \[hep-lat\]](#).
- [4] I. P. Omelyan, I. M. Mryglod and R. Folk, *Comp. Phys. Comm.* **151** (2003) 272.
- [5] Tetsuya Takaishi and Philippe de Forcrand, *Phys. Rev. E* **73** (2006) 036706.
- [6] M. A. Clark, A. D. Kennedy and P. J. Silva, *PoS(LATTICE 2008)041*, [arXiv:0810.1315 \[hep-lat\]](#).
- [7] A. D. Kennedy, M. A. Clark and P. J. Silva, *PoS(LAT2009)021*, [arXiv:0910.2950 \[hep-lat\]](#).
- [8] M.A. Clark, B. Joo, A.D. Kennedy, P.J. Silva,

- PoS(LATTICE2010)323, [arXiv:1011.0230](https://arxiv.org/abs/1011.0230) [[hep-lat](https://arxiv.org/archive/hep)].
- [9] C. Urbach, K. Jansen, A. Shindler and U. Wenger. *Comp. Phys. Comm.* **174** (2005) 87.
- [10] R. G. Edwards and B. Joo, *Nucl. Phys. Proc. Suppl.* **140** (2005) 832.
- [11] in preparation.