

Update on the gradient flow scale on the 2+1+1 HISQ ensembles

**Alexei Bazavov,^{a,b,*} Claude Bernard,^c Carleton E. DeTar,^d Aida X. El-Khadra,^{e,f}
Elvira Gámiz,^g Steven Gottlieb,^h Anthony V. Grebe,ⁱ Urs M. Heller,^j William I. Jay,^k
Andreas S. Kronfeldⁱ and Yin Lin^{k,l}**

^a*Department of Computational Mathematics, Science, Michigan State University, East Lansing, Michigan 48824, USA*

^b*Department of Physics and Astronomy, Michigan State University, East Lansing, Michigan 48824, USA*

^c*Department of Physics, Washington University, St. Louis, Missouri 63130, USA*

^d*Department of Physics and Astronomy, University of Utah, Salt Lake City, Utah 84112, USA*

^e*Department of Physics, University of Illinois, Urbana, Illinois 61801, USA*

^f*Illinois Center for Advanced Studies of the Universe, University of Illinois, Urbana, Illinois 61801, USA*

^g*CAFPE and Departamento de Física Teórica y del Cosmos, Universidad de Granada, E-18071 Granada, Spain*

^h*Department of Physics, Indiana University, Bloomington, Indiana 47405 USA*

ⁱ*Theory Division, Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

^j*American Physical Society, Hauppauge, New York 11788, USA*

^k*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

^l*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions*

We report on the ongoing effort of improving the determination of the gradient flow scale on the (2+1+1)-flavor HISQ ensembles generated by the MILC collaboration. We compute the scales $\sqrt{t_0}/a$ and w_0/a with the Wilson and Symanzik flow using three discretizations for the action density: clover, Wilson and tree-level Symanzik improved. For the absolute scale setting, we intend to employ the Ω -baryon mass, but are also using the pion decay constant while the Ω -mass calculations are in progress.

FERMILAB-CONF-23-0833-T

MIT-CTP/5663

The 40th International Symposium on Lattice Field Theory (Lattice 2023)

July 31st - August 4th, 2023

Fermi National Accelerator Laboratory

*Speaker

1. Introduction

For lattice-QCD precision projects such as the calculation of the hadronic contribution to the anomalous magnetic moment of the muon, a major contribution to the final uncertainty comes from the determination of the lattice spacing (see, for instance, Ref. [1]). A convenient procedure for determining the relative lattice scale with high statistical precision is the gradient flow introduced in Ref. [2]. For the absolute scale, several quantities have been used in the past including the pion or kaon decay constants. However, the Ω -baryon mass, which is measured experimentally to about 0.02% precision, has recently become a more popular choice for precision lattice projects.

The MILC collaboration has been generating a library of (2+1+1)-flavor ensembles with the Highly Improved Staggered Quark (HISQ) action [3] for over a decade now [4, 5]. The available lattice spacings range from 0.15 down to 0.03 fm and the quark masses from approximately 300 MeV down to the physical pion mass. The initial tuning was done with the r_0/r_1 scale and later improved with the f_π scale. In 2015, the w_0 scale was calculated on the existing ensembles [6]. Since then, more ensembles have been added, and some physical mass ensembles were regenerated with more precisely tuned quark masses. Therefore, we are revisiting w_0 -scale setting on the MILC HISQ ensembles. Improvements in this calculation include using a different integrator [7, 8], two types of gradient flow (Wilson and Symanzik), three different discretizations of the observable, and use of the Ω -baryon mass for the absolute scale setting. Recent developments, *e.g.*, Refs. [9–11], allow us to use a fully unitary formalism where the valence quarks are also HISQ. Recently, the CalLat collaboration computed the physical w_0 scale with the Ω -baryon mass input using most of the MILC HISQ publicly available ensembles including the ones with heavier-than-physical pion mass and a few additional HISQ ensembles generated by them [12]. The main difference with our work is that they used the Möbius domain-wall fermion action for the valence quarks and that we compute the Ω -baryon mass only on the physical pion mass ensembles.

2. Gradient flow

2.1 Definitions

The gradient flow [2] is a smoothing procedure where a gauge field configuration is evolved in a fictitious time t according to a diffusion equation,

$$\frac{dV_{x,\mu}}{dt} = - \{ \partial_{x,\mu} S^f(t) \} V_{x,\mu}, \quad V_{x,\mu}(t=0) = U_{x,\mu}. \quad (1)$$

In this project, for the flow action S^f , we use both the Wilson and tree-level improved Symanzik actions. To set the lattice spacing, one evaluates a flowed observable, a popular choice being the action density S^o itself, at flow time $t = t_0$ or $t = w_0^2$ where the corresponding condition is satisfied:

$$t^2 \langle S^o(t) \rangle \Big|_{t=t_0} = 0.3 \quad \text{or} \quad \left[t \frac{d}{dt} t^2 \langle S^o(t) \rangle \right] \Big|_{t=w_0^2} = 0.3. \quad (2)$$

The major practical difference between the resulting $\sqrt{t_0}$ or w_0 scales is higher suppression of the discretization effects in w_0 , as was first noted in Ref. [13]. Moreover, the discretization effects at the tree level that result from a specific combination of the gauge action S^g (with HISQ we

use a one-loop Symanzik-improved action), flow action S^f and observable discretization S^o were calculated in Ref. [14] and can be canceled by replacing the observable with

$$t^2 S^o(t) \rightarrow \frac{t^2 S^o(t)}{1 + \sum_{m=1}^4 C_m (a^{2m}/t^m)}. \quad (3)$$

For the observable, we use Wilson (W), tree-level Symanzik-improved (S) and clover (C) discretization of the action density. In what follows, three letter abbreviations refer to the combination of the gauge, flow, and observable actions, *e.g.*, the SSC (Symanzik-Symanzik-Clover) combination was the only one used in our previous work [6].

2.2 Lattice setup

In general, computation of the gradient flow for a given ensemble requires a smaller computational effort than other typical calculations, *e.g.*, hadron spectrum measurements. However, given the diffusive nature of the flow equation (1) the flow time is actually measured in units of lattice spacing squared. This means that for finer ensembles the flow needs to be run significantly longer (in lattice units) to reach the t_0 or w_0^2 scales. For this reason, we explored a number of integration schemes that would allow the flow integration to have the largest possible step size. We use a fourth-order commutator-free low-storage Lie group integrator described in Refs. [7, 8]. To completely control the systematic errors from the flow integration, we perform all measurements with two step sizes $\Delta t = 1/20$ and $\Delta t = 1/40$. We observe that the integration is stable for all ensembles and the systematic error is always a few orders of magnitude smaller than the statistical error.

We illustrate how different choices influence the discretization effects in Fig. 1 for the 300 MeV pion, $a = 0.12$ fm ensemble. In all panels, we show six different combinations of the flow/observable (with the gauge action being fixed at ‘‘S’’). In the top left panel of Fig. 1, the flows for the $\sqrt{t_0}$ scale are shown and the spread of different curves crossing the 0.3 line hints at relatively large discretization effects. The vertical line shows the scale that corresponds to the SSC combination that was used in our 2015 work [6]. In the top right panel, we show the effect of applying the corrections, Eq. (3), to the observable. The decrease in the spread between different flow/observable combinations shows some suppression of the discretization effects, but they are still quite large for the $\sqrt{t_0}$ scale at $a = 0.12$ fm. Moving to the bottom left panel of Fig. 1, we find that the w_0 scale has drastically reduced discretization effects, and applying the corrections to w_0 , bottom right panel of Fig. 1, has a relatively small effect. Note also that applying the corrections to the derivative (w_0) scale tends to amplify the artifacts at early flow time. The effect of applying the corrections to the $\sqrt{t_0}$ and w_0 scale is similar for the finer ensembles, so we do not show them here.

2.3 Autocorrelations

The scale $\sqrt{t_0}$ or w_0 became a popular choice due to a relatively small computational cost and high statistical precision that can be achieved with either quantity. However, the smearing introduced by the flow, in general, induces longer autocorrelation times for these quantities, compared to r_0 or even f_π scales.

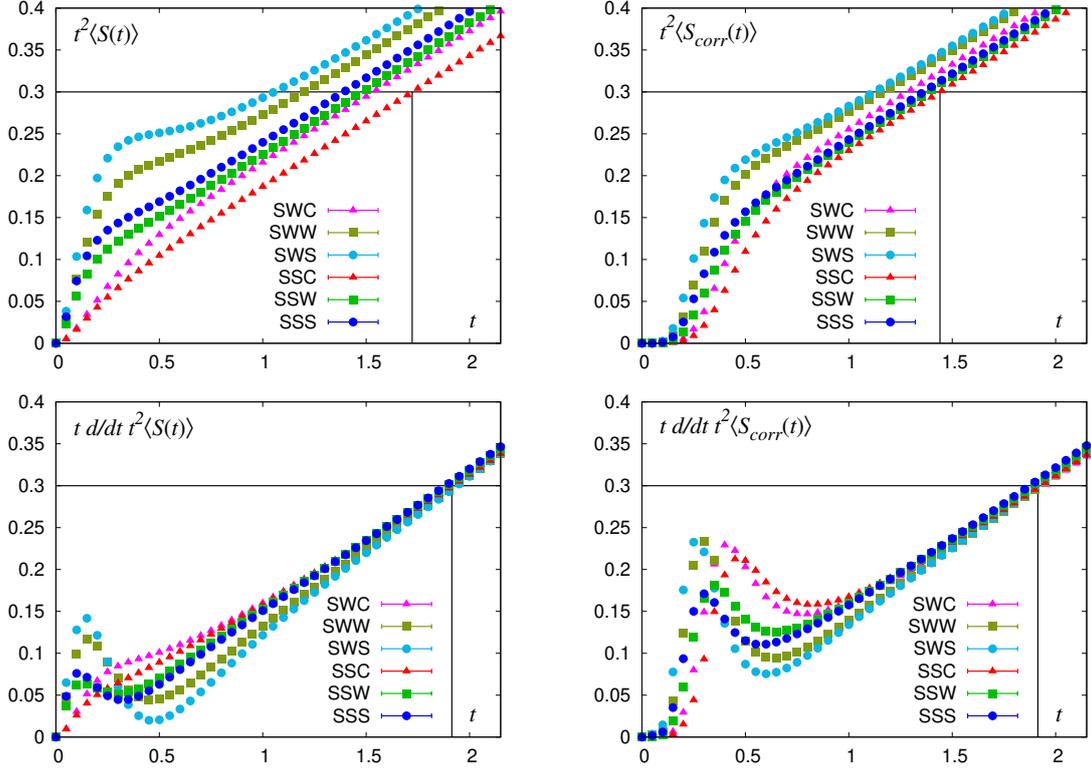


Figure 1: The flow time dependence (in lattice units) of different combinations of the flow and observable. Top panels show the observables used to extract the $\sqrt{t_0}$ scale and the bottom panels the w_0 scale. Left panels show the original observables, while right panels show the effect of applying tree-level corrections for discretization effects, Eq. (3). The statistical errors are smaller than the symbol sizes.

The autocovariance function for an observable O is defined as

$$C(t_{MC}) \equiv \langle (O_i - \langle O_i \rangle)(O_j - \langle O_j \rangle) \rangle = \langle O_0 O_{t_{MC}} \rangle - \langle O \rangle^2, \quad t_{MC} = |i - j|, \quad (4)$$

where i, j index the Monte Carlo time series. The autocorrelation function is simply $c(t_{MC}) \equiv C(t_{MC})/C(0)$. The estimator of the integrated autocorrelation time with the window method is

$$\bar{\tau}_{int}(t_{MC}) = 1 + 2 \sum_{t'=1}^{t_{MC}} c(t'). \quad (5)$$

In general, the quantity (5) is expected to reach a plateau, but in practice this is complicated by the fact that the variance of $\bar{\tau}_{int}(t_{MC})$ diverges. To reliably estimate the integrated autocorrelation time requires significant statistics. To propagate the statistical errors on the autocorrelation function and the integrated autocorrelation time, we use a block jackknife resampling with 20 and 40 blocks to check for the dependence on the block size.

If one assumes that the autocorrelation function is dominated by a single exponential, *i.e.*,

$$c(t_{MC}) = \exp\{-t_{MC}/\tau_{MC}^1\}, \quad (6)$$

the integrated autocorrelation time is given by the sum of the geometric series

$$\tau_{int}^1 = \frac{\exp\{1/\tau_{MC}^1\} + 1}{\exp\{1/\tau_{MC}^1\} - 1}. \quad (7)$$

This provides an alternative way of determining the integrated autocorrelation time by fitting the early time behavior of the autocorrelation function that can be determined with higher statistical precision.

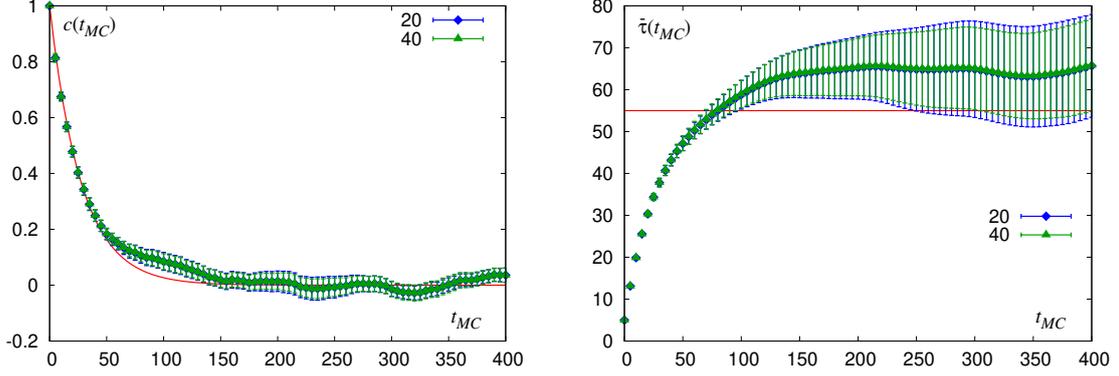


Figure 2: The autocorrelation function (left) and the integrated autocorrelation time (right) measured with the window method for $a = 0.12$ fm ensemble vs. the simulation time in molecular dynamics time units. Different colors correspond to varying the block size in the jackknife procedure.

For the retuned $a = 0.12$ fm physical mass MILC HISQ ensemble, a large data set of 45,000 molecular dynamics time units has been accumulated (with lattices saved every 5 time units). To study autocorrelations we take the clover action density for the observable at time $t = w_0^2$, $O \equiv S^o(t = w_0^2)$, *i.e.*, the quantity from whose ensemble average the w_0 scale is determined. In the left panel of Fig. 2, we show the normalized autocorrelation function $c(t_{MC})$ for that ensemble together with the fit to the form (6) versus the Monte Carlo time in molecular dynamics time units. In the right panel of the figure, we show the integrated autocorrelation time $\bar{\tau}_{int}(t_{MC})$ computed with the window method together with the $\tau_{int}^1 = 55 \pm 3$ computed from assuming a single-exponential decay (7). One can see that from such a large ensemble the autocorrelation function can be determined with enough statistical precision, and as a result one gets a reliable estimate of the integrated autocorrelation time. The value calculated from the single-exponential decay is reasonably consistent with the plateau. Conservatively, we can estimate that the integrated autocorrelation time is below 80 molecular dynamics time units, and the ensemble has at least $45,000/80 \sim 560$ statistically independent samples as far as w_0 is concerned.

Another high-statistics MILC HISQ ensemble at finer lattice spacing is the $a = 0.09$ fm physical-mass ensemble with about 20,000 molecular dynamics time units. Similarly to the previous case, we look at the clover action density at the flow time $t = w_0^2$. Its autocorrelation function is shown in the left panel and the integrated autocorrelation time in the right panel of Fig. 3. The integrated autocorrelation time determined from the single-exponential decay is $\tau_{int}^1 = 43 \pm 3$ molecular dynamics time units and is consistent with the plateau achieved in $\bar{\tau}_{int}(t_{MC})$.

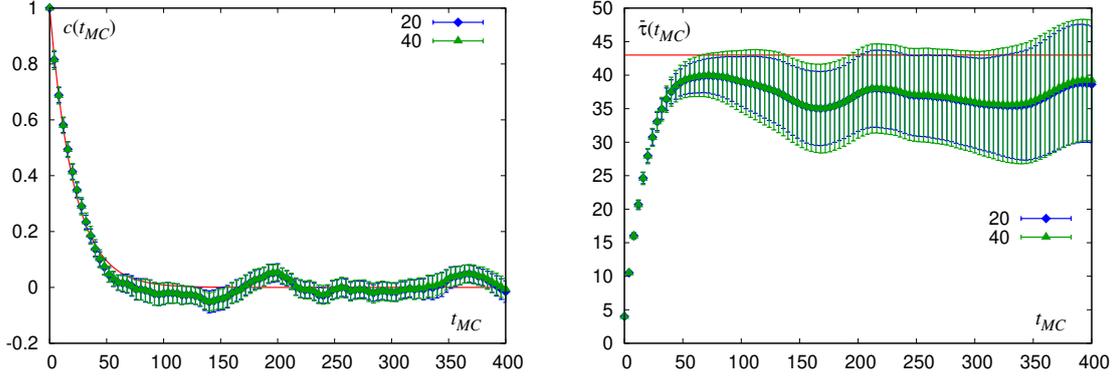


Figure 3: The autocorrelation function (left) and the integrated autocorrelation time (right) measured with the window method for $a = 0.09$ fm ensemble vs. the simulation time in molecular dynamics time units. Different colors correspond to varying the block size in the jackknife procedure.

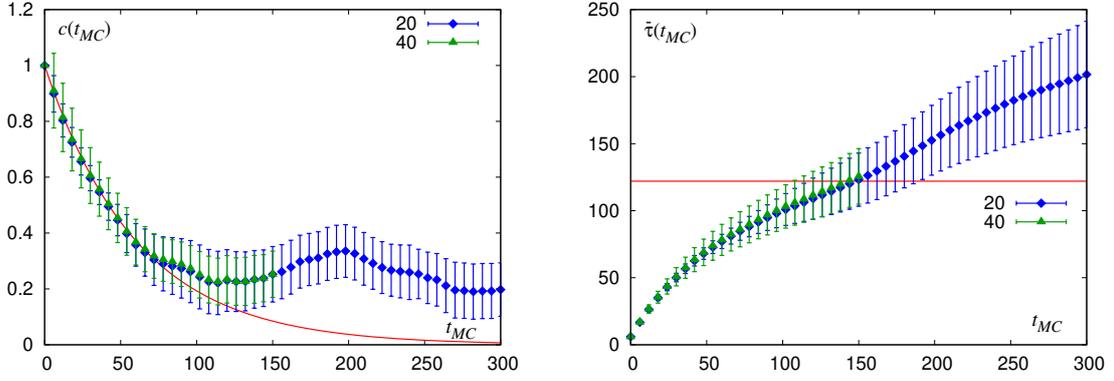


Figure 4: The autocorrelation function (left) and the integrated autocorrelation time (right) estimated with the window method for $a = 0.06$ fm ensemble vs. the simulation time in molecular dynamics time units. Different colors correspond to varying the block size in the jackknife procedure. Note that with 40 blocks the autocorrelation function is computed only up to $t_{MC} = 6,000/40 = 150$ molecular dynamics time units.

We now turn to a finer, $a = 0.06$ fm, 300 MeV pion mass MILC HISQ ensemble that contains 6,000 molecular dynamics time units (1,000 configurations separated by 6 time units). We expect that when decreasing the lattice spacing the integrated autocorrelation time for w_0 related observables will, in general, grow (in molecular dynamics time units). As before, we take the clover action density at $t = w_0^2$ as an observable and show its autocorrelation function in the left panel and the integrated autocorrelation time in the right panel of Fig. 4. Compared to the two previous cases, the errors on the autocorrelation function are much larger and the integrated autocorrelation time does not reach a plateau. It is plausible, based on our experience with the other ensembles, that the increase in the autocorrelation function around 200 time units is a statistical fluctuation. If the Monte Carlo statistics were significantly increased, the autocorrelation function (at least visually) would become more consistent with the single-exponential form, as was observed for the two large-statistics ensembles discussed previously. In that case, the integrated autocorrelation time might

plateau around 120–150 time units, becoming consistent with the estimate $\tau_{int}^1 = 122 \pm 31$ time units. However, at present, the data do not allow one to rule out larger integrated autocorrelation times up to 250 time units. Conservatively, we can estimate that the $a = 0.06$ fm 300 MeV pion MILC HISQ ensemble contains from $6,000/250 = 24$ to $6,000/120 = 50$ statistically independent samples.

Overall, we conclude that for the ensembles with the lattice spacing $a > 0.06$ fm the integrated autocorrelation time for the gradient flow scale observables is on the order of 50 time units or below (it is mainly sensitive to the lattice spacing and less to the values of the quark masses) while for ≤ 0.06 fm ensembles the integrated autocorrelation time may be significant, and one has to exercise care when estimating the true variance of the w_0 scale.

We have currently completed the flow measurements on all existing MILC HISQ ensembles, and the final analysis is ongoing. The details will be presented in a forthcoming publication.

3. Staggered baryons

For the absolute scale setting, we have been computing the Ω -baryon mass on the physical pion mass MILC HISQ ensembles with lattice spacings $a \approx 0.15, 0.12, 0.09,$ and 0.06 fm. On each ensemble, we construct the full set of operators that interpolate to different staggered fermion tastes of Ω baryons as described in Ref. [15]. We use three interpolators here. Two of three operators are mixed and interpolate to the same two taste-split Ω baryons in the ground state; on the other hand, the remaining operator interpolates to a single Ω baryon. In the continuum limit, the taste-splitting mass differences disappear and a single Ω -baryon mass is obtained.

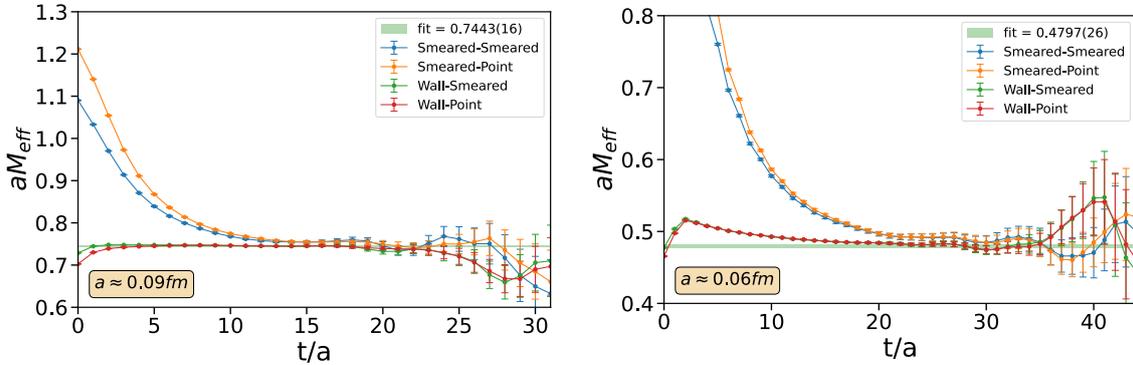


Figure 5: Effective mass plots of the Ω baryon from two-point correlation functions using Gaussian-smeared and corner-wall sources and Gaussian-smeared and point sinks at two lattice spacings ($a = 0.06, 0.09$ fm). A weighted average over a few time slices has been applied to the correlation functions to suppress the effects of oscillations from negative-parity states [11].

For each operator, we use four different source-sink constructions: the corner-wall source (see Eq. (2.1) in Ref. [9] for the definition) with the point and Gaussian-smeared sinks and the Gaussian-smeared source with the point and Gaussian-smeared sink. Even though correlators constructed from corner-wall sources have the best signal-to-noise ratios, the correlators we construct with

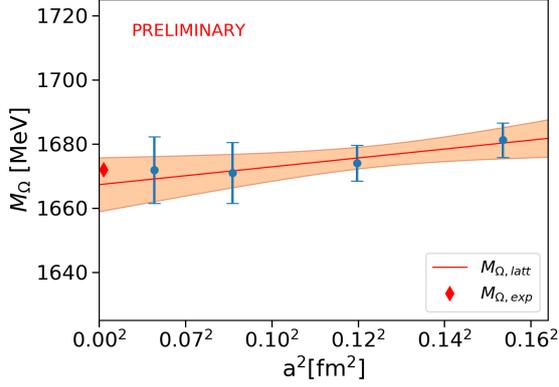


Figure 6: A preliminary continuum extrapolation of the Ω -baryon mass from measurements on ensembles ranging from 0.15 fm down to 0.06 fm using Eq. (8). The lattice spacings are set using f_π from Ref. [16]. The continuum-extrapolated value gives good agreement with the experimental measurement $M_{\Omega, \text{exp}}$.

corner-wall sources are not positive definite because correlators constructed with corner-wall sinks have unfavorable signal-to-noise scaling in volume and cannot be used in practice. So we supplement corner-wall source with Gaussian-smeared sources to constrain the excited state contamination. In preliminary fits, we have analyzed the correlators constructed from the operator that interpolates to a single Ω -baryon taste for simplicity. We also apply a weighted average over a few time slices [11] before fitting to alleviate the excited-state contamination from negative-parity states. We then perform Bayesian fits to all four correlators on each ensemble with two positive-parity states and one negative-parity state to constrain the left-over negative-parity state contribution. Fit results from the two finest lattice spacings are shown in Fig. 5.

Finally, we perform a preliminary continuum extrapolation of the Ω -baryon mass as a consistency check. Here we follow the same procedure used in our earlier work [10] and fit the masses from different lattice spacings, $M_{\text{lat.}, \Omega}(a)$, to the functional form

$$M_{\text{lat.}, \Omega} = M_\Omega (1 + c_1 (a\Lambda_{\text{QCD}})^2 + c_2 (a\Lambda_{\text{QCD}})^4) \quad (8)$$

where we choose $\Lambda_{\text{QCD}} = 200$ MeV, c_1 is unconstrained, the prior on c_2 is $0(1)$, and the lattice spacings are determined using the mass independent scheme with f_π [16]. The result is shown in Fig. 6. The final continuum-extrapolated mass is $M_\Omega = 1668(8)$ MeV, which is consistent with the PDG result within a one sigma uncertainty. We are currently working on increasing statistics on all ensembles to further reduce the uncertainties.

4. Conclusion

The effort to determine the w_0 scale on all MILC HISQ ensembles is ongoing. The combinations of two flows (Wilson and Symanzik) and three observables (Wilson, Symanzik and clover) have been carried out on all existing MILC (2+1+1)-flavor HISQ ensembles and the final error analysis is in progress. The Ω -baryon mass is being computed on the physical pion mass ensembles for the absolute scale setting. Early analysis indicates that an error of well below 0.5% may be achievable.

Acknowledgements

Computations for this work were carried out in part with resources provided by the USQCD Collaboration, the National Energy Research Scientific Computing Center (NERSC) which are funded by the Office of Science of the U.S. Department of Energy as well as resources provided under the ASCR Leadership Computing Challenge (ALCC) program at the Argonne Leadership Computing Facility, a U.S. Department of Energy (DOE) Office of Science user facility at Argonne National Laboratory and is based on research supported by the U.S. DOE Office of Science-Advanced Scientific Computing Research Program, under Contract No. DE-AC02-06CH11357. This work used Ranch tape storage system at Texas Advanced Computing Center (TACC) at The University of Texas at Austin through allocation MCA93S002: Lattice Gauge Theory on Parallel Computers from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program [17], which is supported by National Science Foundation grants Nos. 2138259, 2138286, 2138307, 2137603, and 2138296. Computations on the Big Red II+, Big Red 3 and Big Red 200 supercomputers were supported in part by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute. The parallel file system employed by Big Red II+ was supported by the National Science Foundation under Grant No. CNS-0521433. Small-scale computations utilized computational resources and services provided by the Institute for Cyber-Enabled Research at Michigan State University.

This work was supported in part by the U.S. Department of Energy, Office of Science, under Awards No. DE-SC0010120 (S.G.), No. DE- AC02-07CH11359 (A.V.G and A.S.K), No. DE-SC0011090 (W.I.J. and Y.L.), No. DE-SC0021006 (W.I.J.), No. DE-SC0015655 (A.X.K.), and the Funding Opportunity Announcement Scientific Discovery through Advanced Computing: High Energy Physics, LAB 22-2580; by the National Science Foundation under Grants Nos. PHY20-13064 and PHY23-10571 (C.E.D.), No. PHY18-12332 (A.B.) and No. PHY-2019786 (Y.L.) ; by the Simons Foundation under their Simons Fellows in Theoretical Physics program (A.X.K.); by SRA (Spain) under Grant No. PID2019-106087GB-C21 / 10.13039/501100011033 (E.G.); by the Junta de Andalucía (Spain) under Grants No. FQM-101, A-FQM-467-UGR18 (FEDER), and P18-FR-4314 (E.G.). This document was prepared by the Fermilab Lattice, and MILC Collaborations using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE- AC02-07CH11359.

References

- [1] A. Bazavov *et al.* [Fermilab Lattice, HPQCD, and MILC], *Phys. Rev. D* **107**, 114514 (2023) doi:10.1103/PhysRevD.107.114514 [arXiv:2301.08274 [hep-lat]].
- [2] M. Lüscher, *JHEP* **08**, 071 (2010) [erratum: *JHEP* **03**, 092 (2014)] doi:10.1007/JHEP08(2010)071 [arXiv:1006.4518 [hep-lat]].
- [3] E. Follana *et al.* [HPQCD and UKQCD], *Phys. Rev. D* **75**, 054502 (2007) doi:10.1103/PhysRevD.75.054502 [arXiv:hep-lat/0610092 [hep-lat]].

- [4] A. Bazavov *et al.* [MILC], Phys. Rev. D **82**, 074501 (2010) doi:10.1103/PhysRevD.82.074501 [arXiv:1004.0342 [hep-lat]].
- [5] A. Bazavov *et al.* [MILC], Phys. Rev. D **87**, 054505 (2013) doi:10.1103/PhysRevD.87.054505 [arXiv:1212.4768 [hep-lat]].
- [6] A. Bazavov *et al.* [MILC], Phys. Rev. D **93**, 094510 (2016) doi:10.1103/PhysRevD.93.094510 [arXiv:1503.02769 [hep-lat]].
- [7] A. Bazavov, BIT Num. Math. **62**, 745 (2022) doi:10.1007/s10543-021-00892-x [arXiv:2007.04225 [math.NA]].
- [8] A. Bazavov and T. Chuna, [arXiv:2101.05320 [hep-lat]].
- [9] Y. Lin, A.S. Meyer, C. Hughes, A.S. Kronfeld, J.N. Simone and A. Strelchenko, Phys. Rev. D **103**, 034501 (2021) doi:10.1103/PhysRevD.103.034501 [arXiv:1911.12256 [hep-lat]].
- [10] C. Hughes, Y. Lin and A.S. Meyer, PoS **LATTICE2019**, 057 (2019) doi:10.22323/1.363.0057 [arXiv:1912.00028 [hep-lat]].
- [11] Y. Lin, A.S. Meyer, S. Gottlieb, C. Hughes, A.S. Kronfeld, J.N. Simone and A. Strelchenko, Phys. Rev. D **103**, 054510 (2021) doi:10.1103/PhysRevD.103.054510 [arXiv:2010.10455 [hep-lat]].
- [12] N. Miller, L. Carpenter, E. Berkowitz, C.C. Chang, B. Hörz, D. Howarth, H. Monge-Camacho, E. Rinaldi, D.A. Brantley and C. Körber, *et al.* Phys. Rev. D **103**, 054511 (2021) doi:10.1103/PhysRevD.103.054511 [arXiv:2011.12166 [hep-lat]].
- [13] S. Borsányi *et al.* [BMW], JHEP **09**, 010 (2012) doi:10.1007/JHEP09(2012)010 [arXiv:1203.4469 [hep-lat]].
- [14] Z. Fodor, K. Holland, J. Kuti, S. Mondal, D. Nogradi and C.H. Wong, JHEP **09**, 018 (2014) doi:10.1007/JHEP09(2014)018 [arXiv:1406.0827 [hep-lat]].
- [15] J.A. Bailey, Phys. Rev. D **75**, 114505 (2007) doi:10.1103/PhysRevD.75.114505 [arXiv:hep-lat/0611023 [hep-lat]].
- [16] A. Bazavov, C. Bernard, N. Brown, C.E. Detar, A.X. El-Khadra, E. Gámiz, S. Gottlieb, U.M. Heller, J. Komijani and A.S. Kronfeld, *et al.* Phys. Rev. D **98**, 074512 (2018) doi:10.1103/PhysRevD.98.074512 [arXiv:1712.09262 [hep-lat]].
- [17] T.J. Boerner, S. Deems, T.R. Furlani, S.L. Knuth and J. Towns. 2023. ACCESS: Advancing Innovation: NSF’s Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support. “In Practice and Experience in Advanced Research Computing (PEARC ’23)”, July 23–27, 2023, Portland, OR, USA. ACM, New York, NY, USA. doi.org/10.1145/3569951.3597559.