

Twisted mass gauge ensembles at physical values of the light, strange and charm quark masses

Jacob Finkenrath,^{a,*} Constantia Alexandrou,^{a,b} Simone Bacchio,^a Martha Constantinou,^c Petros Dimopoulos,^d Roberto Frezzotti,^e Karl Jansen,^f Bartosz Kostrzewa,^g Giannis Koutsou,^a Giancarlo Rossi,^{e,h} Carsten Urbachⁱ and Urs Wenger^j

^aComputation-based Science and Technology Research Center, The Cyprus Institute, Nicosia, 20, Constantinou Kavafi Str., Nicosia 2121, Cyprus

^bDepartment of Physics, University of Cyprus, P.O. Box 20537, Nicosia 1678, Cyprus

^cTemple University, 1801 N Broad Str., Philadelphia, PA 19122, USA

^dDipartimento di Scienze Matematiche, Fisiche e Informatiche, Università di Parma and INFN, Gruppo Collegato di Parma, Italy

^eDipartimento di Fisica, Università di Roma "Tor Vergata" and INFN, Sezione di Roma 2, Via della Ricerca Scientifica - 00133 Rome, Italy

^fDeutsches Elektronen-Synchrotron DESY, Platanenallee 6, 15738 Zeuthen, Germany

^gHigh Performance Computing and Analytics Lab, Rheinische Friedrich-Wilhelms-Universität Bonn, Friedrich-Hirzebruch-Allee 8, D-53115 Bonn, Germany

^hCentro Fermi - Museo Storico della Fisica e Centro Studi e Ricerche Enrico Fermi, Compendio del Viminale, Piazza del Viminale 1, I-00184, Rome, Italy

ⁱHISKP (Theory), Rheinische Friedrich-Wilhelms-Universität Bonn, Bonn, Germany

^jAlbert Einstein Center for Fundamental Physics, Institute for Theoretical Physics, University of Bern, Sidlerstrasse 5, CH-3012 Bern, Switzerland

E-mail: j.finkenrath@cyi.ac.cy

Lattice QCD simulations directly at physical masses of dynamical light, strange and charm quarks are highly desirable especially to remove systematic errors due to chiral extrapolations. However such simulations are still challenging. We discuss the adaption of efficient algorithms, like multi-grid methods or higher order integrators, within the molecular dynamic steps of the Hybrid Monte Carlo algorithm, that are enabling simulations of a new set of gauge ensembles by the Extended Twisted Mass collaboration (ETMC). We present the status of the on-going ETMC simulation effort that aims to enabling studies of finite size and discretization effects. We work within the twisted mass discretization which is free of odd-discretization effects at maximal twist and present our tuning procedure.

*The 38th International Symposium on Lattice Field Theory, LATTICE2021 26th-30th July, 2021
Zoom/Gather@Massachusetts Institute of Technology*

*Speaker




Figure 1: Overview of the current $N_f = 2 + 1 + 1$ twisted mass clover-improved ensembles. The left panel shows the ensembles according to their lattice extent in units of the pion mass, while the right panel shows them according to the pion mass. The dark blue circles show the non-physical point ensembles, the light blue already existing physical point ensembles and the red on-going simulations at the physical point.

1. Introduction

During the past decade, progress has been made enabling simulations of twisted mass fermions at physical quark masses at several lattice spacings less than 0.1 fm and volumes as large as 9 fm [1, 2]. These ensembles drive the rich physics program of the Extended Twisted Mass collaboration (ETMC), which ranges from measurements of quark masses, precision measurements of CKM-matrix elements, hadron spectroscopy and scattering, nucleon structure, semi-leptonic decays and many other ETMC projects.

In this presentation, we overview the progress on the algorithmic and computational side that made the generation of these ensembles possible. More specifically, we discuss how to fix the parameters of the simulations, which needs a careful fine tuning procedure in order to achieve $O(a)$ improvement, how to optimize the multigrid solver DDalphaAMG [3, 4] and how to improve the force computation during the Hybrid Monte Carlo simulations. The use of a three level multigrid procedure comes with some limitation for current HPC machines equipped with CPUs that limit scalability. Furthermore, we present selected results using the generated statistics and discuss autocorrelations using the physical point ensembles.

2. Status of Extended Twisted Mass Collaboration $N_f = 2 + 1 + 1$ Simulations

Within ETMC, twisted mass clover-improved ensembles has been generated at four different lattice spacings, namely $a \sim 0.093$ fm, $a \sim 0.08$ fm, $a \sim 0.069$ fm and $a \sim 0.057$ fm, referred to A-, B-, C- and D-lattices, respectively. We target pion mass ranges between 250 MeV and 135 MeV, with the exception of the A-lattices where ensembles with up to 350 MeV are generated. Ensembles with physical pion masses, also denoted as physical point ensembles, are generated at the B-, C- and D-lattice spacings. As depicted in Fig. 1, all physical point ensembles fulfil $m_\pi \cdot L > 3.6$. A larger volume ensemble is generated at the B-lattice spacing with $m_\pi \cdot L = 5.3$, to be used for finite

ensemble	β	c_{SW}	κ	V/a^4	$a\mu_\ell$	$a\mu_\sigma$	$a\mu_\delta$
cA211.53.24	1.726	1.74	0.1400645	$24^3 \times 48$	0.00530	0.1408	0.1521
cA211.40.24				$24^3 \times 48$	0.00400		
cA211.30.32				$32^3 \times 64$	0.00300		
cA211.12.48			0.1400650	$48^3 \times 96$	0.00120		
cB211.25.24	1.778	1.69	0.1394267	$24^3 \times 48$	0.00250	0.1246864	0.131052
cB211.25.32				$32^3 \times 64$	0.00250		
cB211.25.48				$48^3 \times 96$	0.00250		
cB211.14.64			0.1394265	$64^3 \times 128$	0.00140		
cB211.072.64				$64^3 \times 128$	0.00072		
cB211.072.96				$96^3 \times 192$	0.00072		
cC211.20.48	1.836	1.6452	0.13875285	$48^3 \times 96$	0.00060	0.106586	0.107146
cC211.125.64				$64^3 \times 128$	0.00125		
cC211.06.80				$80^3 \times 160$	0.00060		
cD211.17.64	1.900	1.6112	0.137972174	$64^3 \times 128$	0.00170	0.087911	0.086224
cD211.054.96				$96^3 \times 192$	0.00054		

Table 1: Simulation parameters of the $N_f = 2 + 1 + 1$ twisted mass clover-improved ensembles.

volume studies. We target statistics exceeding 3000 MDUs for all ensembles. Parameters are listed in Table 1 and can be also be found in Ref. [1].

2.1 Twisted mass fermion action

The $N_f = 2 + 1 + 1$ ensembles are generated using the Iwasaki gauge action for the pure gauge, the $N_f = 2$ mass-degenerate twisted mass fermion action with a clover term for the light quarks and the $N_f = 1 + 1$ non-degenerate twisted mass fermion action for the strange and charm quarks. The non-degenerated twisted mass operator is given in the heavy doublet flavor space by

$$D(\kappa, c_{SW}, \mu, \epsilon) = D_W(\kappa, c_{SW}) \otimes 1 + i\mu_\sigma \gamma_5 \otimes \tau_3 - \mu_\delta \otimes \tau_1 = \begin{bmatrix} D_W + i\gamma_5 \mu_\sigma & -\mu_\delta \\ -\mu_\delta & D_W - i\gamma_5 \mu_\sigma \end{bmatrix} \quad (1)$$

with D_W the clover improved Wilson Dirac operator, κ the Wilson hopping parameter, c_{SW} the clover parameter and μ_σ and μ_δ the 1+1 twisted mass parameters [5]. Note that taking $\mu_\delta = 0$ the mass-degenerate twisted mass operator in flavor space is recovered. Twisted mass fermions have several advantages. Lattice artefacts of odd power in a can be removed if the Partially Conserved Axial Current (PCAC) mass is tuned to zero, i.e. $m_{PCAC}(\kappa) \rightarrow 0$ [6]. This can be done by fine tuning the bare Wilson quark mass parameter $\bar{m} = 0.5/\kappa - 4$ to its critical mass $\bar{m}_{crit} = 0.5/\kappa_{crit} - 4$. Moreover, in the case of a finite twisted mass value, the operator D is non-singular. For the squared operator one gets in fact

$$D^\dagger D = D_W^\dagger D_W + \mu^2, \quad (2)$$

with μ playing the role of an infra-red cut-off for the eigenvalues of D . While this guarantees convergence of iterative methods, such as the conjugate gradient solver, the twisted mass term breaks isospin symmetry. This results in a mass-splitting within the pion triplet. The neutral pion

mass is shifted with respect to the charged pion mass, which is given in next to leading order (NLO) chiral perturbation theory [7] by

$$a^2(m_{\pi^0}^2 - m_{\pi^\pm}^2) = -c_0 \cdot a^2. \quad (3)$$

This cut-off effect, if large, can give neutral pion mass of zero at finite quark masses and trigger a phase transition. For values of the neutral pion mass close to zero, tuning the PCAC mass to zero becomes notoriously difficult and such Monte Carlo simulations of twisted mass fermions become impossible. This prevented simulation at physical pion masses without a clover term.

Including a clover term, reduces lattice artefacts which decreases the mass gap between the charged and neutral pion given in Eq. (3) to such an extent as to make simulations at values of $a \sim 0.09$ possible [9]. For our setup we use 1-loop tadpole boosted perturbation theory [8] to fix the value of the c_{SW} parameter, given by

$$c_{SW} = 1 + 0.113(3) \frac{6}{\beta \langle P \rangle} \quad (4)$$

with $\langle P \rangle$ the plaquette.

2.2 Parameter tuning

In order to simulate $N_f = 2 + 1 + 1$ twisted mass fermion ensembles at the physical point, we need to set the bare parameters for our action. The complete set of parameters is

$$\{\beta, c_{SW}, \kappa, \mu_\ell, \mu_\sigma, \mu_\delta\}. \quad (5)$$

Using an estimated initial guess, we can pre-select the value of the gauge coupling β that fixes c_{SW} via Eq. (4). The bare-mass parameters of the twisted mass action that are left, namely

$$\{\kappa, \mu_\ell, \mu_\sigma, \mu_\delta\} \quad (6)$$

require careful fine tuning in order to guarantee $O(a)$ improvement. This can be done by tuning the bare Wilson quark mass κ towards its critical value by requiring

$$\frac{Z_A m_{PCAC}(\kappa, \mu_\ell, \mu_\sigma, \mu_\delta)}{\mu_\ell} < 0.1, \quad (7)$$

with Z_A the axial renormalization factor. Note that the PCAC mass depends also on the heavy quark parameters μ_σ and μ_δ of the non-degenerated twisted mass operator. Due to this dependence on the heavy quark pair, charm and strange, we utilise Osterwalder Seiler (OS) fermions [31] to set the strange and charm quark mass. We then match the non-unitary setup of the OS-fermions with the unitary setup with non-degenerate twisted mass fermions. This results in three tuning conditions for the heavy quark parameters. The first two are given by

$$C_1 = \frac{\mu_c^{OS}}{\mu_s^{OS}} = 11.8 \quad \text{and} \quad C_2 = \frac{m_{D_s}^{OS}}{f_{D_s}^{OS}} = 7.9 \quad (8)$$

that set the strange and charm quark mass parameters in the OS sector. The full matching between OS and non-degenerate twisted fermions is finally done using the kaon mass m_K by requiring

$$m_K^{OS}(\mu_s^{OS}) \equiv m_K^{ND}(\mu_\sigma, \mu_\delta), \quad (9)$$

which basically utilizes the renormalization factor ratio Z_P/Z_S and yields to the determination of μ_σ and μ_δ via the relation $\mu_{c,s}^{OS} = \frac{1}{Z_P} \left(\mu_\sigma \pm \frac{Z_P}{Z_S} \mu_\delta \right)$.




Figure 2: The tuning procedure employed for the D-lattice spacings using reweighting on two tuning ensembles. We depict the dependence of the PCAC mass on the light twisted mass parameter $a\mu_\ell$ (left), on the Wilson bare quark mass $a\bar{m}$ (middle) the heavy quark parameter μ_σ (right).

3. Parameter tuning

For the tuning of the twisted mass parameters, we iterate the following steps: i) tune κ at constant $\{\mu_\ell, \mu_\sigma, \mu_\delta\}$; ii) change μ_ℓ and retune κ ; iii) tune $\{\mu_\sigma, \mu_\delta\}$ and retune κ ; iv) reiterate until stability is reached. This procedure requires for the tuning of the A-, B- and C- ensembles to generate for each steps roughly two Markov chains with around 500 MDU, see [2] for a detailed discussion for the case of the B-lattice spacing. For the D-lattice we modify the tuning procedure by making use of reweighting for all parameters $\{\kappa, \mu_\ell, \mu_\sigma, \mu_\delta\}$ [10–12]. This reduces the number of ensembles needed for the tuning, namely to one per iteration circle.

With the help of reweighting by changing the bare parameters by 1% to 5% we could give an estimate of the first derivatives of $am_{PCAC}(\kappa, \mu_\ell, \mu_\sigma, \mu_\delta)$. This quantity was estimated in the case of the tuning procedure of the A-, B- and C-lattice spacings via generation of several ensembles.

For the tuning of the D-lattice spacing we generate two different ensembles at twisted mass values $a\mu = 0.002$ and $a\mu = 0.00125$ and volumes of $V = 32^3 \times 64$ and $V = 48^3 \times 96$, respectively. We employ the fit Ansatz

$$f(\bar{m}, \mu_\ell, \mu_\ell) = c_0 + c_1 \bar{m} + c_2 \mu_\ell + c_3 \mu_\ell \quad (10)$$

with bare quark mass $\bar{m} = 1/2/\kappa - 4$. We find for the fit coefficients $c_0 = 0.51(2)$, $c_1 = 1.37(6)$, $c_2 = 0.04(2)$ and $c_3 = 0.012(4)$. The resulting uncorrelated fit with $\chi^2 = 11.2$ is shown in Fig. 2, where we include all available data points generated by reweighting resulting in 19 degrees of freedom. The critical mass parameter at physical light quark mass of $a\mu = 0.00054$ is given by $\kappa_{crit}^{D,(0)} = 0.137973465$. Based on this estimate we generate 500 MDUs at our the physical point with $a\mu = 0.00054$, leading to a slightly negative PCAC mass of $-4.4(7)10^{-5}$. This does meet our criterion and thus we proceed to retune κ using our estimate for the slope $\partial f / \partial \bar{m} = 1.37$ such that our final estimate for the critical Wilson mass parameter is given by $\kappa_{crit}|_{\beta=1.9} = 0.137972174$.

With this tuning procedure, that it is also outlined in ref. [2] for the B-lattice spacing, we were able to successfully tune towards critical Wilson mass, achieving $O(a)$ -improvement. Indeed as seen in Fig. 3, all ensembles fulfil the condition of eq. 7 with $m_{PCAC}/\mu < 4\%$ or better.




Figure 3: The history of the PCAC mass for the physical point ensembles is plotted in units of molecular dynamics (MDUs) showing from top to bottom the cB211.072.64, cB211.072.96, cC211.060.80 and the two replicas of cD211.060.80 ensembles.

4. Simulation setup

For the generation of the ensembles listed in Table 1, we use our open source software suite *tmLQCD* [14–16], which implements an optimised Hybrid Monte Carlo algorithm enabling the use of twisted mass operators of Eq. (1) and improved gauge action with rectangular loops. For the molecular dynamics, we are using, for the mass-degenerated light quark doublet, even-odd Hasenbusch mass preconditioning with masses $\{\rho_0 = \mu_\ell, \rho_1, \dots, \rho_N\}$, while for the heavy quark doublet, even-odd rational approximation of the square root of the non-degenerate twisted mass operator $\hat{Q}_{ND}^2 = D_{ND, eo} D_{ND, eo}^\dagger$ is used [13]. Thus, the Boltzmann weight of our setup is given by

$$W(U) = Z^{-1} \exp \left\{ -\beta S_{iwa}(U) + \text{Tr} \ln \{W_{oo}(\rho_0, \mu_\sigma)\} - \sum_{j=1}^N \phi^\dagger \left[\frac{q_i}{\hat{Q}_{ND}^2 + \mu_i} \right] \phi \right. \\ \left. - \prod_{j=1}^{N-1} \eta_j^\dagger (1 + \Delta^2 \rho_{j-1, j} (\hat{Q}^2 + \rho_{j-1}^2)^{-1} \eta_j - \eta_N^\dagger (\hat{Q}^2 + \rho_N^2)^{-1} \eta_N \right\}.$$

For a more detailed discussion see Appendix A of Ref. [1]. Our simulation code *tmLQCD* provides currently a link to the algebraic multigrid solver library DDalphaAMG, which provides routines for the twisted mass operator [4] and for the non-degenerated twisted mass operator [18] that can speed-up the smallest shifts within the rational approximation. Additionally, we have employed the mixed-precision linear solver provided by the software package *QPhiX* [19–23], which can be utilized for the larger mass shifts, where the algebraic multigrid solver becomes less effective.




Figure 4: The requirements of a three level multigrid method within HMC simulations are shown. The left panel shows the scalability for three different volumes on SuperMUC-NG. The middle panel shows the memory requirements and the right panel illustrates the reversibility violation as a function of the solver residual.

4.1 Application of HMC with multigrid solver

For all simulations of the physical point ensembles, we take advantage of the highly reduced computational costs of using multigrid solvers within the force calculation, see e.g. [4]. The use of multigrid solver within the integration of Hamiltons molecular dynamics requires some additional care and comes with some limitations. Namely, the scalability on HPC systems, such as SuperMUC-NG, is limited by the volume of the coarsest grid within the multigrid procedure. This limits the strong scaling window, as depicted in Fig. 4, breaking down for a three-level MG method for a lattice of size $V = 64^3 \times 128$ at around 80 Skylake nodes, and for a lattice of size $V = 96^3 \times 192$ at around 420 Skylake nodes. This results in a roughly scaling of the upper bound of the strong scaling window by L^3 or $V^{3/4}$ and limits the maximal effective parallelisation of our HMC.

The lower bound of the working window using a multigrid solver is determined by memory requirements. In fact prolongation and restriction operators required for projecting iteration vectors from level to level and building up the coarser operators, need allocation of $\mathcal{O}(20)$ full vectors, a number which scales with the size of the volume. On SuperMUC-NG with 192 GB RAM per node this introduces a hard limit for the minimal parallelisation, given by about 16 nodes for lattice size of $V = 64^3 \times 128$ and increasing to 80 nodes for a volume of $V = 96^3 \times 192$.

The usage of a multigrid solver not only limits the scalability window, but can also compromise the correctness of the HMC sampling. By reusing and updating the coarse grid operators and the corresponding prolongation and restriction operators from previous integration steps, the reversibility criterium, which is needed to fulfil detailed balance, is violated. In order to assess the magnitude of the effect, we studied reversibility violation within the HMC using a test volume with size $V = 32^3 \times 64$ and measuring the variance of $\delta\Delta H$ for different solvers. $\delta\Delta H$ is given by the difference of the Hamiltonian at the beginning of the trajectory with the Hamiltonian integrated to $\tau = 1$ and integrated back to the start. As discussed in Ref. [17] for variances below 0.01 no deviations from the expectation values, such as the plaquette, are found. To match the same precision in $\sigma^2(\delta\Delta H)$ as the mixed-precision solver of QPhiX, the square residual of the multigrid solver needs to be two orders of magnitudes smaller. For all our physical point ensembles, we checked reversibility and find that for the selected solver criteria (see [1]), $\sigma^2(\delta\Delta H)$ is below the bound of set in Ref. [17].




Figure 5: Comparison of different numerical integrators. The variance of the energy violation is plotted as a function of the step size for a test ensemble of volume $V = 24^3 \times 48$.

4.2 Nested force gradient integrator

The computational cost of molecular dynamics scales with the number of integration steps. At constant acceptance rate it follows $\text{cost} \propto V^{1+\frac{1}{2n}}$ with n the order of the selected integrator. Thus higher order schemes have a better volume scaling. A class of various different symplectic reversible integration schemes are discussed in Ref. [26]. This includes schemes improved with force gradient terms. For example, the second minimal norm scheme can be extended to fourth order

$$\Delta(h) = e^{h\frac{1}{6}\hat{B}} e^{h\frac{1}{2}\hat{A}} e^{h\frac{2}{3}\hat{B}-\frac{1}{12}h^3\hat{C}} e^{h\frac{1}{2}\hat{A}} e^{h\frac{1}{6}\hat{B}} \quad (11)$$

by including the force gradient term C , which is given by $C = 2 \sum_{x=1}^V \sum_{\nu=0}^3 \frac{\partial S}{\partial U_{\nu}(x)} \frac{\partial^2 S}{\partial U_{\nu}(x) \partial U_{\mu}(x)}$. The additional second derivate term, needed in the force gradient term, can be approximated by an additional force term, as outline in Ref. [25]. This not only reduces the cost of the calculation of one force calculation but also simplifies the application by implicitly taking care of cross-terms between different parts of the actions. It turns out that the force gradient improves the minimal norm scheme of Eq. 11 and outperforms for larger volumes the other integrators, as shown in Fig. 5 for a $L = 24$ lattice. Note that to tune the nested integrator setup we minimise the cost function at constant acceptance rate using an effective model for the higher order terms in line with Ref. [27].

4.3 Computation costs

A summary of the improvements and the achieved reduction in computational costs per HMC trajectory are shown in Fig. 6. By enabling DDalphaAMG within the HMC, there is a reduction of the computational cost per trajectory by more than one order of magnitude [4]. In addition, the adaptation of DDalphaAMG to the non-degenerate twisted mass operator speeds up the non-degenerated sector, as discussed in Ref. [18]. The use of higher order integrators has given further improvements, especially going to larger volumes, such as $L = 96$ [28]. Note that there is still space for further improvements, for example, by coarse level improvements, which are currently under investigation. Moreover, a multi right handside version of DDalphaAMG is available, which, however to utilize, would require larger refraction of the force computation [29]. We are investigating how this potentially could be adapted within a lattice QCD python API [32], which is currently under development.




Figure 6: The computational cost of a trajectory for $N_f = 2 + 1 + 1$ simulations as a function of the lattice spatial extent L .

4.4 Autocorrelations at the physical point

One major unsolved challenge in lattice QCD with periodic boundary conditions is adequate sampling of different topological sectors at very fine lattice spacings [30]. We have monitored the gauge definition of the topological charge Q at gradient flow time t_0 . As expected for our range of lattice spacings between $a \sim 0.057$ fm to $a \sim 0.082$ fm, the topological charge is fluctuating well between topological sectors. Moving towards finer lattice spacing, we have currently indication for an increase of the autocorrelation time, hinting at the fact that for simulations below $a < 0.045$ fm further algorithmic improvements will be needed in order to sample the topological charge properly within 3000 MDUs.

5. Conclusions

With the current ensembles we are able to study finite volume and lattice spacing artefacts directly at the physical point. Using the three ensembles cB211.072.64, cC211.060.80 and cD211.054.96, we can take the continuum limit, and with the two ensembles at the B-lattice spacing we can study finite volume effects. In the future, we are planning to simulate at larger volumes and explore new approaches for enabling simulations at lattice spacings smaller than 0.05 fm, which are currently limited by the critical slowing down of the algorithms and their scalability with the volume. Due to the behavior of the strong scaling window of our multigrid solver DDalphaAMG, the real time per trajectories will further increase resulting in longer generation times per ensemble. We are taking this issue into account in our future software development, by improving and enabling linking to state-of-the-art QCD software libraries within tmLQCD, such as QUDA, as well as by developing a new user-friendly flexible python API *lynxs* [32]. This will enable HMC simulations on the next generation of high performance systems.




Figure 7: The history of the topological charge Q at gradient flow time t_0 of the physical point runs plotted versus the molecular dynamics (MDUs). We show from top to bottom cB211.072.64, cB211.072.96, cC211.060.80 and two different replicas for cD211.060.80.

6. Acknowledgments

We thank all members of the ETM collaboration for a most conducive cooperation. J.F. and S.B. are supported by the H2020 project PRACE 6-IP (GA No. 82376) and the EuroCC (GA No. 951740). P. D. acknowledges support from the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 813942 (EuroPLEx) and from INFN under the research project INFN-QCDLAT. G.K. acknowledges support from project NextQCD, co-funded by the European Regional Development Fund and the Republic of Cyprus through the Research and Innovation Foundation (RIF) (EXCELLENCE/0918/0129). Partial support is provided by the H2020 European Joint Doctorate program STIMULATE grant No. 765048. We acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for project pr74yo by providing computing time on SuperMUC at LRZ (www.lrz.de). The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this proceeding. This work used resources from NIC on JUWELS and Jureca Booster at the JSC, under projects with ids ECY00, HBN28 and HCH02. We acknowledge PRACE for awarding us access to HAWK at HLRS, where part of our work is carried out within the project with Id Acid 4886.

References

- [1] C. Alexandrou *et al.* [Extended Twisted Mass], Phys. Rev. D **104**, no.7, 074520 (2021) doi:10.1103/PhysRevD.104.074520 [arXiv:2104.06747 [hep-lat]].
- [2] C. Alexandrou, S. Bacchio, P. Charalambous, P. Dimopoulos, J. Finkenrath, R. Frezzotti,

- K. Hadjiyiannakou, K. Jansen, G. Koutsou and B. Kostrzewa, *et al.* Phys. Rev. D **98**, no.5, 054518 (2018) doi:10.1103/PhysRevD.98.054518 [arXiv:1807.00495 [hep-lat]].
- [3] A. Frommer, K. Kahl, S. Krieg, B. Leder and M. Rottmann, SIAM J. Sci. Comput. **36**, A1581-A1608 (2014) doi:10.1137/130919507 [arXiv:1303.1377 [hep-lat]].
- [4] C. Alexandrou, S. Bacchio, J. Finkenrath, A. Frommer, K. Kahl and M. Rottmann, Phys. Rev. D **94**, no.11, 114509 (2016) doi:10.1103/PhysRevD.94.114509 [arXiv:1610.02370 [hep-lat]].
- [5] R. Frezzotti and G. C. Rossi, Nucl. Phys. B Proc. Suppl. **128**, 193-202 (2004) doi:10.1016/S0920-5632(03)02477-0 [arXiv:hep-lat/0311008 [hep-lat]].
- [6] R. Frezzotti and G. C. Rossi, JHEP **08**, 007 (2004) doi:10.1088/1126-6708/2004/08/007 [arXiv:hep-lat/0306014 [hep-lat]].
- [7] S. R. Sharpe and J. M. S. Wu, Phys. Rev. D **71**, 074501 (2005) doi:10.1103/PhysRevD.71.074501 [arXiv:hep-lat/0411021 [hep-lat]].
- [8] S. Aoki, R. Frezzotti and P. Weisz, Nucl. Phys. B Proc. Suppl. **73**, 915-917 (1999) doi:10.1016/S0920-5632(99)85243-8 [arXiv:hep-lat/9809179 [hep-lat]].
- [9] A. Abdel-Rehim *et al.* [ETM], Phys. Rev. D **95**, no.9, 094515 (2017) doi:10.1103/PhysRevD.95.094515 [arXiv:1507.05068 [hep-lat]].
- [10] A. Hasenfratz, R. Hoffmann and S. Schaefer, Phys. Rev. D **78**, 014515 (2008) doi:10.1103/PhysRevD.78.014515 [arXiv:0805.2369 [hep-lat]].
- [11] J. Finkenrath, F. Knechtli and B. Leder, Nucl. Phys. B **877**, 441-456 (2013) [erratum: Nucl. Phys. B **880**, 574-575 (2014)] doi:10.1016/j.nuclphysb.2013.10.019 [arXiv:1306.3962 [hep-lat]].
- [12] C. Alexandrou, J. Finkenrath, L. Funcke, K. Jansen, B. Kostrzewa, F. Pittler and C. Urbach, Phys. Rev. Lett. **125**, no.23, 232001 (2020) doi:10.1103/PhysRevLett.125.232001 [arXiv:2002.07802 [hep-lat]].
- [13] C. Urbach, K. Jansen, A. Shindler and U. Wenger, Comput. Phys. Commun. **174**, 87-98 (2006) doi:10.1016/j.cpc.2005.08.006 [arXiv:hep-lat/0506011 [hep-lat]].
- [14] K. Jansen and C. Urbach, Comput. Phys. Commun. **180**, 2717-2738 (2009) doi:10.1016/j.cpc.2009.05.016 [arXiv:0905.3331 [hep-lat]].
- [15] A. Abdel-Rehim, F. Burger, A. Deuzeman, K. Jansen, B. Kostrzewa, L. Scorzato and C. Urbach, PoS LATTICE2013, 414 (2014) doi:10.22323/1.187.0414 [arXiv:1311.5495 [hep-lat]].
- [16] A. Deuzeman, K. Jansen, B. Kostrzewa and C. Urbach, PoS LATTICE2013, 416 (2014) doi:10.22323/1.187.0416 [arXiv:1311.4521 [hep-lat]].
- [17] C. Urbach, Comput. Phys. Commun. **224**, 44-51 (2018) doi:10.1016/j.cpc.2017.12.005 [arXiv:1710.07526 [hep-lat]].

- [18] C. Alexandrou, S. Bacchio and J. Finkenrath, *Comput. Phys. Commun.* **236**, 51-64 (2019) doi:10.1016/j.cpc.2018.10.013 [arXiv:1805.09584 [hep-lat]].
- [19] B. Joó, D. D. Kalamkar, K. Vaidyanathan, M. Smelyanskiy, K. Pamnany, V. W. Lee, P. Dubey and W. Watson, *Lect. Notes Comput. Sci.* **7905**, 40-54 (2013) doi:10.1007/978-3-642-38750-0_4
- [20] Joó, Bálint and Kalamkar, Dhiraj D and Kurth, Thorsten and Vaidyanathan, Karthikeyan and Walden, Aaron International Conference on High Performance Computing, 415–427, 2016
- [21] Joó, Bálint and Smelyanskiy, Mikhail and Kalamkar, Dhiraj D and Vaidyanathan, Karthikeyan techreport, 2015
- [22] S. Heybrock, B. Joó, D. D. Kalamkar, M. Smelyanskiy, K. Vaidyanathan, T. Wettig and P. Dubey, doi:10.1109/SC.2014.11 [arXiv:1412.2629 [hep-lat]].
- [23] Joó, B. and Kalamkar, D.D. and Vaidyanathan, K. and Smelyanskiy, M. and Kurth, T. and Walden, A. and Schroeck, M. and Labus, P. and Ueding, M. and Kostrzewa, B. and Dep, D., <https://github.com/JeffersonLab/qphix>
- [24] A. D. Kennedy, P. J. Silva and M. A. Clark, *Phys. Rev. D* **87**, no.3, 034511 (2013) doi:10.1103/PhysRevD.87.034511 [arXiv:1210.6600 [hep-lat]].
- [25] H. Yin and R. D. Mawhinney, *PoS LATTICE2011*, 051 (2011) doi:10.22323/1.139.0051 [arXiv:1111.5059 [hep-lat]].
- [26] I. P. Omelyan, I. M. Mryglod, and R. Folk, *Comput.Phys.Commun.* 151 272
- [27] M. A. Clark, B. Joo, A. D. Kennedy and P. J. Silva, *PoS LATTICE2010*, 323 (2010) doi:10.22323/1.105.0323 [arXiv:1011.0230 [hep-lat]].
- [28] D. Shcherbakov, M. Ehrhardt, J. Finkenrath, M. Günther, F. Knechtli and M. Pardon, *Commun. Comput. Phys.* **21**, no.4, 1141-1153 (2017) doi:10.4208/cicp.OA-2016-0048 [arXiv:1512.03812 [math.NA]].
- [29] S. Yamamoto, S. Bacchio and J. Finkenrath, *PoS (LATTICE2021)* 536
- [30] S. Schaefer *et al.* [ALPHA], *Nucl. Phys. B* **845**, 93-119 (2011) doi:10.1016/j.nuclphysb.2010.11.020 [arXiv:1009.5228 [hep-lat]].
- [31] R. Frezzotti and G. C. Rossi, *JHEP* **10**, 070 (2004) doi:10.1088/1126-6708/2004/10/070 [arXiv:hep-lat/0407002 [hep-lat]].
- [32] S. Bacchio, J. Finkenrath and C. Stylianou, *PoS (LATTICE2021)* 542