

# Chiral rank- $k$ truncations for the multigrid preconditioner of Wilson fermions in lattice QCD

---

Travis Whyte,<sup>a,\*</sup> Andreas Stathopoulos<sup>b</sup> and Eloy Romero<sup>c</sup>

<sup>a</sup>*Jülich Supercomputing Center,  
Jülich, Germany*

<sup>b</sup>*Department of Computer Science, William & Mary,  
Williamsburg, VA, USA*

<sup>c</sup>*Thomas Jefferson National Accelerator Facility,  
Newport News, VA, USA*

*E-mail:* [t.whyte@fz-juelich.de](mailto:t.whyte@fz-juelich.de)

We present a modification to the setup algorithm for the multigrid preconditioner of Wilson fermions in lattice QCD. A larger number of test vectors than that used in conventional multigrid is generated by the smoother. This set of test vectors is then truncated by a singular value decomposition on the chiral components of the test vectors, which are subsequently used to form the prolongation and restriction matrices of the multigrid hierarchy. This modification is demonstrated to improve the convergence of linear equations on an anisotropic lattice with  $m_\pi \approx 239$  MeV from the Hadron Spectrum Collaboration and an isotropic lattice with  $m_\pi \approx 220$  MeV from the MILC Collaboration. The lattice volume dependence of the method is also examined.

*The 41st International Symposium on Lattice Field Theory (LATTICE2024)  
28 July - 3 August 2024  
Liverpool, UK*

---

\*Speaker

## 1. Introduction

The most efficient and widely used method of solving linear equations for the generation of propagators in lattice QCD is multigrid [1–9]. The multigrid preconditioner is created by exposing the near null space of the Wilson-Dirac operator,  $\mathbf{D}$ , through the generation of test vectors. These test vectors are rich in components that correspond to the low eigenmodes of  $\mathbf{D}$  and are used to create the coarse grid matrices defining the multigrid preconditioner hierarchy. The number of test vectors used to create the coarse grid matrices must be chosen judiciously. Too many test vectors results in larger and denser coarse grid matrices which are expensive to apply. Too few test vectors will result in a preconditioner which cannot effectively reduce the error, leading to an increase in the number of iterations required of the outer iterative solver.

In this article, a modification of the setup method for the multigrid preconditioner of Wilson type fermions is presented, which improves the basis of test vectors used to create the coarse grid matrices by calculating an enlarged basis of test vectors. This basis is subsequently truncated using a singular value decomposition, which retains the components of the test vectors that have the largest contribution to the basis. Numerical results are reported for anisotropic lattices from the Hadron Spectrum Collaboration and isotropic lattices from the MILC Collaboration.

## 2. Multigrid for Wilson Fermions

In the adaptive multigrid framework, the smoother is applied on the homogenous equation

$$\mathbf{D}\psi_i \approx 0, \quad i = 1, \dots, m \quad (1)$$

with  $m$  random initial guesses to create a set of  $m$  near null or test vectors,  $\Psi = [\psi_1, \dots, \psi_m]$  [24]. This reveals the components of the error that the smoother is not able to easily resolve, e.g. the low frequency components of the error. These smooth test vectors are then blocked according to a domain decomposition of the underlying lattice, which then become the columns of the *prolongator* matrix,  $\mathbf{P}$ . The *restriction* matrix,  $\mathbf{R}$ , is typically taken to be  $\mathbf{R} = \mathbf{P}^\dagger$ . The test vectors generated display the phenomenon known as *local coherence*, wherein the smooth modes approximate the low eigenspace of  $\mathbf{D}$  on a given local domain [13, 14]. However, the success of the multigrid preconditioner for the Wilson-Dirac operator relies on “chiral splitting”, wherein the smooth test vectors are split into chiral components using the projectors  $1 \pm \gamma_5$  [1]. The application of these projectors to the smooth test vectors splits the spin components of  $\Psi$  such that

$$\begin{aligned} \Psi^{(\alpha=0)} &= [\psi_1^{(\beta=1,2)}, \dots, \psi_m^{(\beta=1,2)}] \\ \Psi^{(\alpha=1)} &= [\psi_1^{(\beta=3,4)}, \dots, \psi_m^{(\beta=3,4)}], \end{aligned} \quad (2)$$

where  $\alpha$  labels the chiral index and  $\beta$  labels the spin index of the test vectors. The chirally split test vectors are then orthonormalized on a given domain of the spacetime lattice  $\Lambda_j \in \Lambda = \{\Lambda_1, \dots, \Lambda_d\}$  and the prolongator is formed as

$$P_{aij}(x) = \begin{cases} \psi_i^\alpha(x) & \text{if } x \in \Lambda_j \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

where  $x$  denotes the lattice coordinate and color indices have been suppressed. This results in a block diagonal matrix of dimension  $n \times 2md$ , where  $d = |\Lambda|$ . The use of the projectors  $1 \pm \gamma_5$  enforces the symmetry  $\gamma_5 \mathbf{P} = \mathbf{P} \sigma_3$ , where  $\sigma_3$  is the third Pauli spin matrix, which ensures that if a right singular vector is in the range of  $\mathbf{P}$ , a left singular vector is in the range of  $\mathbf{R}$ . The coarse grid matrix is formed by

$$\mathbf{D}_{\ell+1} = \mathbf{R} \mathbf{D}_\ell \mathbf{P}. \quad (4)$$

This process is repeated recursively beginning with the coarse grid operator  $\mathbf{D}_{\ell+1}$  to form successively coarser grids. The projector  $(1 \pm \gamma_5^\ell)$  is used to perform the chiral splitting, where  $\gamma_5^\ell = \sigma_3$  for  $\ell > 0$ . Typical applications of multigrid in lattice QCD utilize the multigrid hierarchy in a  $K$ -cycle [25] as a preconditioner for a flexible outer solver such as FGMRES or GCR.

### 3. Chiral Rank- $k$ Truncations of the Multigrid Basis

The use of rank- $k$  truncations for the multigrid basis of smooth test vectors was given in Ref. [22] and was shown to increase the performance of the multigrid preconditioner for two-dimensional anisotropic heat diffusion equations. It was also used in Ref. [23] to form composite aggregate preconditioners in bootstrap AMG. The extension of this method to Wilson fermions in lattice QCD is straightforward. Guided by the required chiral splitting in the conventional set up method, the singular value decomposition is performed separately on the chiral components of the initial set of  $m$  smooth test vectors, with  $m \geq k$ , given by Eq. (2). Let  $\mathbf{U} \mathbf{\Sigma} \mathbf{V}^\dagger = \mathbf{\Psi}^{(\alpha)}(\Lambda_j)$  be the singular value decomposition of  $\mathbf{\Psi}^{(\alpha)}(\Lambda_j)$ . Then

$$\mathbf{U}_k^{(\alpha)} \mathbf{\Sigma}_k^{(\alpha)} \mathbf{V}_k^{(\alpha)\dagger} \approx \mathbf{P}_{\alpha j} \mathbf{\Psi}^{(\alpha, \ell+1)}(x_c) \quad (5)$$

where  $\alpha = 0, 1$ , is the best rank- $k$  approximation to the initial basis  $\mathbf{\Psi}^{(\alpha)}(\Lambda_j)$ . By matching variables in Eq. (5), we observe that  $\mathbf{U}_k^{(\alpha)} = \mathbf{P}_{\alpha j}$  and  $\mathbf{\Sigma}_k^{(\alpha)} \mathbf{V}_k^{(\alpha)\dagger} = \mathbf{\Psi}^{(\alpha, \ell+1)}(x_c)$ , with  $x_c$  a coarse grid coordinate. The prolongator is then formed as

$$P_{\alpha ij}(x) = \begin{cases} u_i^{(\alpha)}(x) & \text{if } x \in \Lambda_j \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where  $\mathbf{P}$  is of dimension  $n \times 2kd$ . As the singular value decomposition creates a basis of  $m$  orthonormal left singular vectors, the orthonormalization step of the conventional method on the truncated basis is not needed.


### 4. Numerical Experiments

The setup method with chiral rank- $k$  truncations is tested against conventional multigrid for two different types of lattice ensembles. The first, which we refer to as Ensemble A, corresponds to an anisotropic lattice of dimension  $32^3 \times 256$  with a pion of mass  $m_\pi \approx 239$  MeV from the Hadron Spectrum Collaboration [17, 18]. The gauge links arising from these configurations are stout smeared [26]. See Ref. [19] for more information relating to these lattices, including scale and mass determination. The second, which we refer to as Ensemble B, corresponds to an isotropic lattice of dimension  $32^3 \times 64$  with a pion of mass  $m_\pi \approx 220$  MeV from the MILC Collaboration [20, 21]. In

this Clover on HISQ action, the gauge links are smeared with one iteration of HYP smearing [16]. Ref. [15] contains more information on the determination of the relevant physical properties of this ensemble. All experiments were performed on eight nodes with Intel Xeon Gold 6130 processors with  $2 \times 16$  cores. In all experiments, a hierarchy of three multigrid operators is used and the results are averaged over ten right hand sides.

#### 4.1 Size of the Multigrid Basis

As the size of the initial multigrid basis  $m$  is increased, it is natural to expect that more information about the near null space is captured by the preconditioner, leading to improved convergence of the linear equations. Typical applications of multigrid in lattice QCD use between 24 and 32 test vectors. However, the chiral rank- $k$  truncation allows us to calculate an initial basis size of  $m \geq k$  and incorporate information from the test vectors that the conventional setup method ignores. It is thus important to quantify the performance increase offered by increasing  $m$ . Fig. 1




**Figure 1:** The mean solve time as a function of the initial basis size for Ensemble A (left) and Ensemble B (right). The bottom and top y-axis displays  $m$  for levels  $\ell = 0, 1$ , respectively.


shows the mean execution time for solving the systems of linear equations for both Ensemble A and Ensemble B. The basis is truncated at  $k = 24, 32$  for levels  $\ell = 0, 1$ , respectively. A general trend is observed for both ensembles: as  $m$  is increased, the mean execution time decreases. This is inline with expectations that additional information from the exposed smooth modes is made available to the preconditioner. It is also observed that a saturation in the speedup begins between  $m = 96$  and  $m = 128$ , indicating that the modes beyond  $m = 96$  contribute very little to reducing the error.

#### 4.2 Optimal Truncation

The number of test vectors used to create the prolongation and restriction matrices is an important determination. Too few vectors results in a poor preconditioner and too many test vectors results in large and dense coarse grid matrices that are expensive to apply. It is thus important to examine the rank of the truncation utilized. The singular spectra of the test vectors provides some information for this determination. Fig. 2 shows the singular spectrum of the chirally split test vectors for the first domain, which is representative of other domains. It is observed that the spectrum for both ensembles displays a rapid decay of the singular values. The magnitude of the 24th singular value is approximately two orders of magnitude lower than that of the first. This indicates that including left singular vectors into the basis beyond this approximate cutoff will have a small overall contribution to the low rank approximation of the multigrid basis.



**Figure 2:** The singular spectrum for the chirally split test vectors on the first domain for level  $\ell = 0$  (top) and  $\ell = 1$  (bottom) of Ensemble A (left) and Ensemble B (right).



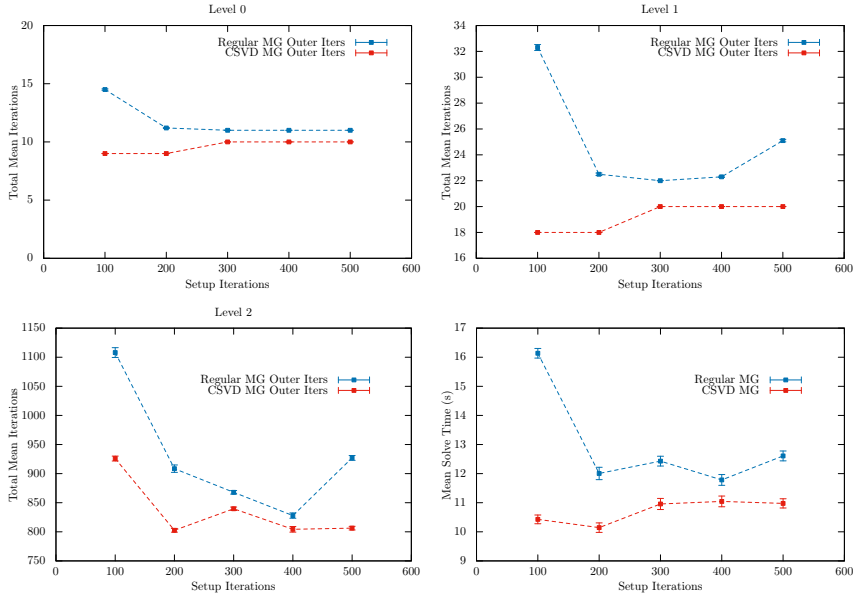
**Figure 3:** The mean execution time for the system of linear equations when the degree of the truncation is varied using an initial basis size of  $m = 96$  for Ensemble A (left) and Ensemble B (right). The bottom and top y-axis displays  $k$  for levels  $\ell = 0, 1$ , respectively.

Fig. 3 displays the mean execution time of the system of linear equations when the degree of the truncation,  $k$ , is varied for a constant value of  $m = 96$ . This value of  $m$  was chosen as larger values of  $m$  provide only a small decrease in execution time. For both ensembles, the optimal truncation is observed to be at  $k = 24$ . For Ensemble A this is comparable to  $k = 16$ , while Ensemble B exhibits a distinct minimum at  $k = 24$ .

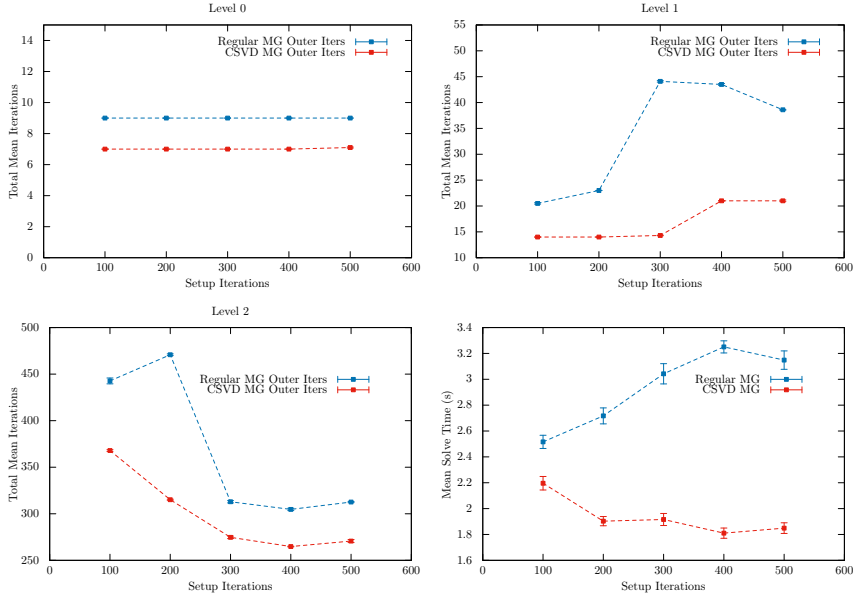
### 4.3 Optimal Number of Setup Iterations

The number of set up iterations of the smoother to generate the test vectors may also have a large effect on the efficacy of the preconditioner. It is thus beneficial to examine the performance of the preconditioner while varying the number of setup iterations while generating the test vectors.

Figs. 4 and 5 display the total number of iterations required on each level  $\ell$  as well as the mean



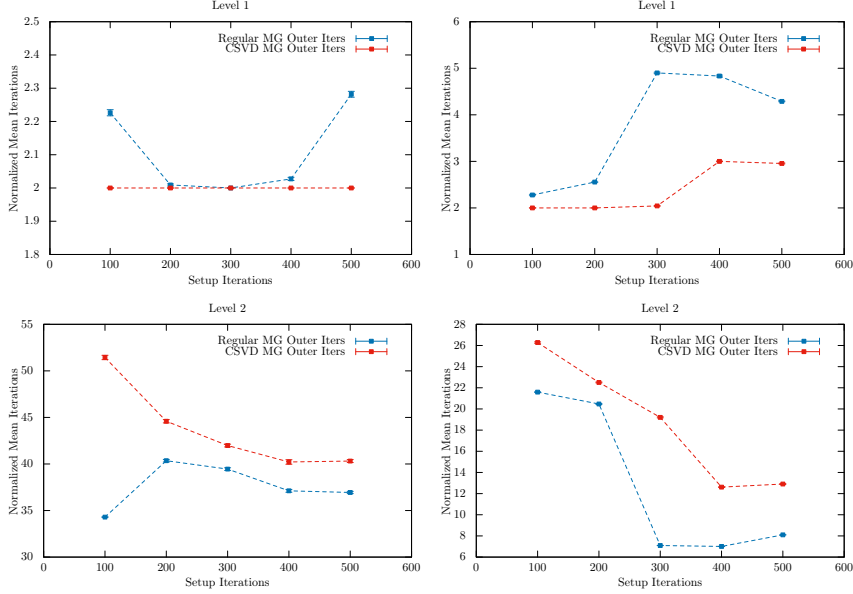
**Figure 4:** The total number of iterations on  $\ell = 0$  (upper left),  $\ell = 1$  (upper right),  $\ell = 2$  (lower left) and the mean solve time of the system of linear equations for Ensemble A.



**Figure 5:** As Fig. 4 for Ensemble B.

execution time. For both ensembles, a value of  $m = 96$  and  $k = 24$  is used, which was found to be the optimal values as observed in Figs. 1 and 3. In both cases, it is observed that the total cost of solving the system of linear equations is reduced in comparison to conventional multigrid for all number of setup iterations on every level. Additionally, the use of the chiral rank- $k$  truncation

results in a preconditioner that is less sensitive to the number of setup iterations.




**Figure 6:** The normalized number of iterations on level  $\ell = 1$  (top) and  $\ell = 2$  (bottom) for Ensemble A (left) and Ensemble B (right).

While the chiral rank- $k$  truncation setup method is observed to decrease the total cost of solving the linear equations, it is important to quantify the performance of the preconditioner on each level per iteration due to the use of the  $K$ -cycle. In a  $K$ -cycle, multiple iterations of the solver on level  $\ell + 1$  occur for each iteration on level  $\ell$ . To examine the cost per iteration on each level, the number of iterations on level  $\ell + 1$  is normalized by the number of iterations on level  $\ell$ . Fig. 6 displays the normalized cost for levels  $\ell = 1, 2$  on both ensembles. It is observed for both ensembles across all number of setup iterations that the cost per iteration is greater for the preconditioner that utilizes the chiral rank- $k$  truncation. This is consistent with expectations that more information about the low frequency components of the error is being transferred to the coarse grids where more iterations can be performed at significantly less cost.

#### 4.4 Lattice Volume Scaling

We now examine the volume dependence of the set up methods with and without the chiral rank- $k$  truncation as the volume of the hypercubic lattice is increased. The system of linear equations are solved on three lattice volumes of spatial extent  $L_S = 24, 32, 40$  and temporal extent  $L_t = 64$  and  $m_\pi \approx 220$  MeV with the Clover on HISQ action<sup>1</sup>. The parameters of the calculation were chosen from the near optimal setup parameters and used across all three volumes. To examine the scaling with the lattice volume, the mean execution time of the system of linear equations preconditioned by both methods is measured across all three volumes, shown in Fig. 7. For all three volumes, the preconditioner using chiral rank- $k$  truncations shows a decrease in execution time in comparison to the conventional method.

<sup>1</sup>The same lattice of  $L_S = 32$  is used here as in the preceding numerical results.



**Figure 7:** The performance of both MG preconditioners as the lattice volume is increased.

## 5. Summary

We have presented a modification to the setup algorithm for the multigrid preconditioner for systems of linear equations of Wilson fermions. The chiral rank- $k$  truncation utilizes a singular value decomposition of the chiral components of the test vectors restricted to a domain of the lattice. By calculating out a large basis of test vectors, the singular value decomposition is able to truncate the basis to the fewest vectors containing the largest contribution to the low rank approximation of the basis. In all numerical experiments, the use of the chiral rank- $k$  truncation results in a decrease in the time required to solve the system of linear equations.

## Acknowledgments

TW acknowledges partial funding from the Exascale Computing Project (ECP), Project Number: 17-SC-20-SC, a collaborative effort of the U.S. Department of Energy, Office of Science and the National Nuclear Security Administration, as well as from a Royal Society Research Fellowship and the U.K. Science and Technology Facilities Council (STFC) [grant numbers ST/T000694/1, ST/X000664/1]. TW also acknowledges support from Science Foundation Ireland [grant number 21/FFP-P/10186] and Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the CRC 1639 NuMerIQS – project no. 511713970. AS and ER acknowledge partial support by DOE SciDAC-5 grant (DE-FOA-0002589) and AS acknowledges partial support from National Science Foundation grant IIS-2008557. Part of this work was performed using computing resources at William & Mary with the software codes Chroma [10], QPhiX [11] and MG\_PROTO [12]. We thank the MILC Collaboration and the Hadron Spectrum Collaboration ([www.hadspec.org](http://www.hadspec.org)) for making the gauge configurations available to us. The gauge configurations of the Hadron Spectrum Collaboration were generated using resources awarded from the U.S. Department of Energy INCITE program at Oak Ridge National Lab, the NSF Teragrid at the Texas Advanced Computer Center and the Pittsburgh Supercomputer Center, as well as at Jefferson Lab. We also thank Steven Gottlieb and Walter Wilcox for their help in obtaining the MILC configurations.



## References

- [1] J. Brannick et al., *Adaptive Multigrid Algorithm for Lattice QCD*, *Phys. Rev. Lett.* **100** (2008) 041601 [0707.4018].
- [2] Richard C. Brower et al., *Multigrid for chiral lattice fermions: Domain wall*, *Phys. Rev. D* **102** (2020) 094517 [2004.07732].
- [3] Richard C. Brower et al., *Multigrid algorithm for staggered lattice fermions*, *Phys. Rev. D* **97** (2018) 114513 [1801.07823].
- [4] James Brannick et al., *Multigrid Preconditioning for the Overlap Operator in Lattice QCD*, *Numer. Math.* **132** (2016) 463 [1410.7170].
- [5] A. Frommer et al., *Adaptive Aggregation-Based Domain Decomposition Multigrid for the Lattice Wilson–Dirac Operator*, *SIAM J. Sci. Comput.* **36** (2014) A1581 [1303.1377].
- [6] S. Bacchio et al., *DDalphaAMG for Twisted Mass Fermions*, *PoS LATTICE2016* (2016) 259 [1611.01034].
- [7] J. Brannick and K. Kahl, *Bootstrap Algebraic Multigrid for the 2D Wilson Dirac system*, *SIAM J. Sci. Comput.* **36** (2014) B321 [1308.5992].
- [8] Travis Whyte and Walter Wilcox and Ronald B. Morgan, *Deflated GMRES with Multigrid for Lattice QCD*, *Phys. Lett. B* **803** (2020) 135281 [1912.02868].
- [9] Ronald B. Morgan et. al., *Two-Grid Deflated Krylov Methods for Linear Equations*, 2005.03070.
- [10] Robert G. Edwards and Balint Joo, *The Chroma software system for lattice QCD*, *Nucl. Phys. Proc. Suppl.* **140** (2005) 832 [hep-lat/0409003].
- [11] Balint Joo et al., *Lattice QCD on Intel® Xeon Phi™ Coprocessors*, *Supercomputing, Lecture Notes in Computer Science*, 7905 40-54 2013 Springer Berlin Heidelberg
- [12] Balint Joo, *MG\_PROTO: A Multigrid Library for QCD*.
- [13] Martin Lüscher, *Deflation acceleration of lattice QCD simulations*, *JHEP* **12** (2007) 011 [0710.5417].
- [14] Martin Lüscher, *Local coherence and deflation of the low quark modes in lattice QCD*. *JHEP* **07** (2007) 081 [0706.2298].
- [15] Rajan Gupta et al., *Isvector Charges of the Nucleon from 2+1+1-flavor Lattice QCD*, *Phys. Rev. D* **98** (2018) 034503 [1806.09006].
- [16] Anna Hasenfratz and Francesco Knechtli, *Flavor symmetry and the static potential with hypercubic blocking*, *Phys. Rev. D* **64** (2001) 034504 [hep-lat/0103029].

- [17] Robert G. Edwards and Balint Joo and Huey-Wen Lin, *Tuning for Three-flavors of Anisotropic Clover Fermions with Stout-link Smearing*, *Phys. Rev. D* 78 (2008) 054501 [0803.3960].
- [18] Huey-Wen Lin et al., *First results from 2+1 dynamical quark flavors on an anisotropic lattice: Light-hadron spectroscopy and setting the strange-quark mass*, *Phys. Rev. D* 79 (2009) 034502 [0810.3588].
- [19] David J. Wilson et al., *The quark-mass dependence of elastic  $\pi K$  scattering from QCD*, *Phys. Rev. Lett.* 123 (2019) 042002 [1904.03188].
- [20] A. Bazovov et al., *Scaling studies of QCD with the dynamical HISQ action*, *Phys. Rev. D* 82 (2010) 074501 [1004.0342].
- [21] A. Bazovov et al., *Lattice QCD Ensembles with Four Flavors of Highly Improved Staggered Quarks*, *Phys. Rev. D* 87 (2013) 054505 [1212.4768].
- [22] Edmond Chow, *An Aggregation Multilevel Method Using Smooth Error Vectors*, *SIAM J. Sci. Comput.* 27 (2006) 1727.
- [23] Pasqua D'Ambra and Panayot S. Vassilevski, *Improving solve time of aggregation-based adaptive AMG*, *Num. Lin. Alg. with App.* 26 (2019) 2269.
- [24] M. Brezina et al., *Adaptive Smoothed Aggregation ( $\alpha$ SA)*, *SIAM J. Sci. Comput.* 25 (2004) 1896.
- [25] Yvan Notay and Panayot S. Vassilevski, *Recursive Krylov-based multigrid cycles*, *Num. Lin. Alg. with App.* 15 (2008) 473.
- [26] Colin Morningstar and Mike Peardon, *Analytic smearing of SU(3) link variables in lattice QCD*, *Phys. Rev. D* 69 (2004) 054501 [hep-lat/0311018].