

Bergische Universität Wuppertal

Fachbereich Mathematik und Naturwissenschaften

Institute of Mathematical Modelling, Analysis and Computational
Mathematics (IMACM)

Preprint BUW-IMACM 20/15

Andreas Frommer, Karsten Kahl, Francesco Knechtli, Matthias
Rottmann, Artur Strebel, Ian Zwaan

**A multigrid accelerated eigensolver for the
Hermitian Wilson-Dirac operator in lattice QCD**

April 28, 2020

<http://www.math.uni-wuppertal.de>

A multigrid accelerated eigensolver for the Hermitian Wilson-Dirac operator in lattice QCD

Andreas Frommer, Karsten Kahl, Francesco Knechtli, Matthias Rottmann,
Artur Strebel, Ian Zwaan

Abstract

Eigenvalues of the Hermitian Wilson-Dirac operator are of special interest in several lattice QCD simulations, e.g., for noise reduction when evaluating all-to-all propagators. In this paper we present a Davidson-type eigensolver that utilizes the structural properties of the Hermitian Wilson-Dirac operator Q to compute eigenpairs of this operator corresponding to small eigenvalues. The main idea is to exploit a synergy between the (outer) eigensolver and its (inner) iterative scheme which solves shifted linear systems. This is achieved by adapting the multigrid DD- α AMG algorithm to a solver for shifted systems involving the Hermitian Wilson-Dirac operator. We demonstrate that updating the coarse grid operator using eigenvector information obtained in the course of the generalized Davidson method is crucial to achieve good performance when calculating many eigenpairs, as our study of the local coherence shows. We compare our method with the commonly used software-packages PARPACK and PRIMME in numerical tests, where we are able to achieve significant improvements, with speed-ups of up to one order of magnitude and a near-linear scaling with respect to the number of eigenvalues. For illustration we compare the distribution of the small eigenvalues of Q on a 64×32^3 lattice with what is predicted by the Banks-Casher relation in the infinite volume limit.

Keywords: lattice QCD, eigensolver, multigrid, Wilson-Dirac operator, parallel computing

1. Introduction

In lattice Quantum Chromodynamics (QCD) the Wilson-Dirac operator describes the interaction between quarks and gluons in the framework of quantum field theory. Results of lattice QCD simulations represent essential input to several of the current and planned experiments in elementary particle physics (e.g., BELLE II, LHCb, EIC, PANDA, BES III).

Obtaining eigenpairs (eigenvalues and -vectors) of the Wilson-Dirac operator is an important computational task. For example, eigenpairs can be used to directly compute physical observables [5, 12, 15, 16, 21, 29] or as a tool for noise reduction in stochastically estimated quantities like disconnected fermion

loops [2]. In most circumstances we are interested in a small to moderate amount of eigenvectors corresponding to the eigenvalues closest to zero, especially for the Hermitian Wilson-Dirac operator. As the Hermitian Wilson-Dirac operator is indefinite, these eigenvalues lie in the *interior* of the spectrum.

Typically, computing interior eigenvalues is particularly expensive, which is why in this paper we develop efficient computational methods for the special case of the Hermitian Wilson-Dirac operator. The most prominent methods for obtaining interior eigenvalues are *shift-and-invert* algorithms which extend the basic inverse iteration approach; cf. [34]. This includes methods ranging from the classical Rayleigh quotient iteration (RQI) [34] to the generalized Davidson (GD) methods [34] and its numerous variations like GD+ k [39], Jacobi-Davidson (JD) [36] or JDCG/JDQMR [30, 39]. The generalized Davidson methods can, alternatively, also be regarded as a generalization of Arnoldi’s method [28]¹ with improved search directions.

Most Davidson-type methods share an *inner-outer-scheme*, where the outer iteration finds approximations to the sought eigenpairs, while the inner iterations generates new search directions by approximately solving shifted linear systems

$$(A - \tau I)x = b, \tag{1.1}$$

where τ is an approximation to a target eigenvalue. In fact, these inversions make up the bulk of the computational work, and it is thus mandatory to find particularly efficient methods for this task. In lattice QCD, adaptive algebraic multigrid methods have established themselves as the most efficient methods for solving linear systems with the Wilson-Dirac operator [1, 6, 18, 19, 31]. They demonstrate significant speed-ups compared to conventional Krylov subspace methods, achieving orders of magnitudes faster convergence and an insensitivity to conditioning. In this work, we use the adaptive domain decomposition algebraic multigrid method DD- α AMG [18, 19], but we expect the results for DD- α AMG to carry over to the other aggregation-based multigrid implementations as well. Originally DD- α AMG is composed of an adaptive aggregation based multigrid construction and a red-black multiplicative Schwarz smoother (traditionally termed “SAP” for Schwarz alternating procedure in lattice QCD), but we also have the option to use other smoothers like GMRES in the DD- α AMG framework.

So far, multigrid solvers have been limited to the unshifted Wilson-Dirac operator². This is due to the algebraic construction of the interpolation operator, which is built from approximations of eigenvectors corresponding to small eigenvalues, which is necessary for an effective overall error reduction in the unshifted case, cf. [18, 19].

During the progress of the eigensolver—when the shift τ in (1.1) becomes

¹In [28] the authors relate generalized Davidson with the Lanczos method, but the statement also holds for non-Hermitian matrices.

²The original solver can manage small shifts, but quickly deteriorates for increasing shifts as will be demonstrated later.

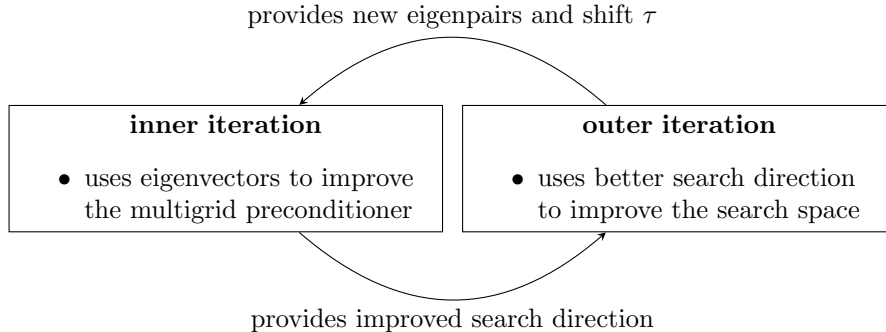


Figure 1: Interleaving eigenpair extraction and construction of the multigrid preconditioner for the inner iteration

larger—the interpolation operator no longer approximates the space spanned by eigenvectors corresponding to small eigenvalues of the shifted operator, thus invalidating the coarse grid correction step and significantly slowing down convergence. The main idea of this paper is that in order to overcome this problem we dynamically update the interpolation operator of the multigrid solver during the outer iteration. This idea is sketched in Figure 1. It illustrates how we are interleaving the eigenpair extraction with the construction of an efficient solver for (1.1). While the outer iteration extracts eigenpairs, its eigenvectors are used to improve the multigrid preconditioner. In turn the multigrid method is an efficient preconditioner for (1.1), which produces improved search directions for the outer iteration.

As an additional topic, we investigate several approaches for the most suitable smoothing method of the multigrid method for the Hermitian Wilson-Dirac operator in presence of large shifts. Altogether, we obtain a method which scales close to linearly with the number of eigenpairs to be computed.

The remainder of this paper is organized as follows. We give a brief introduction into lattice QCD and algebraic multigrid methods in Section 2. We put particular emphasis on showing and theoretically justifying how the algebraic multigrid approach, which has originally been developed for the non-Hermitian Wilson-Dirac operator, can be applied to the Hermitian Wilson-Dirac operator. In Section 3 we proceed by introducing the generalized Davidson method and present our adaptations for the specific case of the Hermitian Wilson-Dirac operator, including the interleaving of the eigensolver and the construction of the multigrid solver used in the inner iteration. Numerical tests and comparisons with commonly used and state-of-the-art software are presented in Section 4. As a simple example of the use of our method in physics, Section 5 discusses lattice artifacts observed for the spectral gap. Finally, a summary of our results is given in Section 6.

2. Algebraic Multigrid Methods in Lattice QCD

The Dirac equation

$$\mathcal{D}\psi + m \cdot \psi = \eta \quad (2.1)$$

describes the dynamics of quarks and the interaction of quarks and gluons. Here, $\psi = \psi(x)$ and $\eta = \eta(x)$ represent quark fields. They depend on x , the points in space-time, $x = (x_0, x_1, x_2, x_3)$. The gluons are represented in the Dirac operator \mathcal{D} , and m is a scalar mass parameter that is independent of x and sets the mass of the quarks in the QCD theory.

More precisely, \mathcal{D} is given as

$$\mathcal{D} = \sum_{\mu=0}^3 \gamma_{\mu} \otimes (\partial_{\mu} + A_{\mu}) , \quad (2.2)$$

where $\partial_{\mu} = \partial/\partial x_{\mu}$ and A is the gluon (background) gauge field with the anti-Hermitian traceless matrices $A_{\mu}(x)$ being elements of $\mathfrak{su}(3)$, the Lie algebra of the special unitary group $SU(3)$. The (Hermitian) γ -matrices $\gamma_0, \gamma_1, \gamma_2, \gamma_3 \in \mathbb{C}^{4 \times 4}$ represent generators of the Clifford algebra with

$$\gamma_{\mu}\gamma_{\nu} + \gamma_{\nu}\gamma_{\mu} = 2\delta_{\mu\nu}I_4 \quad \text{for } \mu, \nu = 0, 1, 2, 3, \quad (2.3)$$

with I_4 the identity on \mathbb{C}^4 . Consequently, at each point x in space-time, the spinor $\psi(x)$, i.e., the quark field ψ at a given point x , is a twelve component column vector, each component corresponding to one of three colors (acted upon by $A_{\mu}(x)$) and four spins (acted upon by γ_{μ}). For future use we remark that $\gamma_5 = \gamma_0\gamma_1\gamma_2\gamma_3$ satisfies

$$\gamma_5\gamma_{\mu} = -\gamma_{\mu}\gamma_5, \quad \mu = 0, 1, 2, 3, \quad (2.4)$$

independently from the chosen representation.

The only known way to obtain predictions in QCD from first principles and non-perturbatively, is to discretize and then simulate on a computer. The discretization is typically formulated on an equispaced periodic $N_t \times N_s^3$ lattice \mathcal{L} with uniform lattice spacing a , N_s denoting the number of lattice points for each of the three spatial dimensions and N_t the number of lattice points in the time dimension. A quark field ψ is now represented by its values at each lattice point, i.e., it is a spinor valued function $\psi : x \in \mathcal{L} \rightarrow \psi(x) \in \mathbb{C}^{12}$.

The *Wilson-Dirac* discretization is one of the most commonly used discretizations in lattice QCD simulations. It is obtained from the continuum equation by replacing the covariant derivatives by centralized covariant finite differences on the lattice. It contains an additional, second order finite difference stabilization term, as otherwise the discretization would suffer from ‘red-black’ instability, cf. [37]. The Wilson-Dirac discretization yields a local operator D in the sense that it represents a nearest neighbor coupling on the lattice \mathcal{L} .

Introducing shift vectors $\hat{\mu} = (\hat{\mu}_0, \hat{\mu}_1, \hat{\mu}_2, \hat{\mu}_3)^T \in \mathbb{R}^4$ in dimension μ on \mathcal{L} , i.e.,

$$\hat{\mu}_{\nu} = \begin{cases} a & \mu = \nu \\ 0 & \text{else} \end{cases} ,$$

the action of D on a discrete quark field ψ is given as

$$(D\psi)(x) = (m_0 + \frac{4}{a})\psi(x) - \frac{1}{2a} \sum_{\mu=0}^3 ((I_4 - \gamma_\mu) \otimes U_\mu(x)) \psi(x + \hat{\mu}) - \frac{1}{2a} \sum_{\mu=0}^3 ((I_4 + \gamma_\mu) \otimes U_\mu^H(x - \hat{\mu})) \psi(x - \hat{\mu}). \quad (2.5)$$

Here, the gauge-links $U_\mu(x)$ are now matrices from the Lie group $SU(3)$, and the lattice indices $x \pm \hat{\mu}$ are to be understood periodically. The mass parameter m_0 sets the quark mass (for further details, see [27]), and we will write $D(m_0)$ whenever the dependence on m_0 is important.

To explicitly describe D we fix a representation for the γ -matrices in which $\gamma_5 = \begin{pmatrix} I_2 & 0 \\ 0 & -I_2 \end{pmatrix}$.

From (2.5) we obtain the couplings of the lattice sites x and $x \pm \hat{\mu}$ as

$$(D)_{x,x+\hat{\mu}} = \frac{-1}{2a}(I_4 - \gamma_\mu) \otimes U_\mu(x), \quad (D)_{x,x-\hat{\mu}} = \frac{-1}{2a}(I_4 + \gamma_\mu) \otimes U_\mu^H(x - \hat{\mu}), \quad (2.6)$$

which shows $(D)_{x+\hat{\mu},x} = \frac{-1}{2a}(I_4 + \gamma_\mu) \otimes U_\mu^H(x)$. Thus, the commutativity relations (2.4) imply the symmetry

$$(\gamma_5 \otimes I_3)(D)_{x,x+\hat{\mu}} = ((\gamma_5 \otimes I_3)(D)_{x+\hat{\mu},x})^H.$$

With $\Gamma_5 = I_{n_{\mathcal{L}}} \otimes \gamma_5 \otimes I_3$, $n_{\mathcal{L}}$ the number of lattice sites, this symmetry can be described on the level of the entire Wilson-Dirac operator as

$$\Gamma_5 D = (\Gamma_5 D)^H. \quad (2.7)$$

The matrix Γ_5 is Hermitian and unitary, and the Γ_5 -symmetry (2.7) is a non-trivial, fundamental symmetry that the discrete Wilson-Dirac operator inherits from a corresponding symmetry of the continuum Dirac operator (2.2).

The Wilson-Dirac operator and its clover-improved variant (where a term which is diagonal in space and time is added to reduce the local discretization error from $\mathcal{O}(a)$ to $\mathcal{O}(a^2)$) is an adequate discretization for the numerical computation of many physical observables. For further details on discretization, its properties and the clover-improved variant, we refer the interested reader to [11, 20, 35]. In this paper we focus on the *Hermitian* or *symmetrized Wilson-Dirac operator*

$$Q := \Gamma_5 D.$$

Algebraic multigrid methods.. The state-of-the-art approaches for solving linear systems involving the (non-Hermitian) Wilson-Dirac operator D are variants of aggregation-based adaptive algebraic multigrid methods, see [1, 6, 19, 31]. Typically, these multigrid solvers are used as a (non-stationary) preconditioner within a flexible Krylov subspace method like FGMRES [33] or GCR [14]; see also [25, 26].

The error propagator for the two-level version of all these multigrid approaches is

$$E_{2g} = (I - MD)^\nu (I - PD_c^{-1}RD)(I - MD)^\mu, \quad (2.8)$$

where M denotes the smoother—which, in the case of DD- α AMG, is given by the Schwarz alternating procedure (SAP)—and μ and ν denote the number of pre- and post-smoothing iterations, respectively. The operator $I - PD_c^{-1}RD$ is the coarse grid correction, where P is the adaptively constructed aggregation based interpolation [1, 6, 19, 31], obtained in a “setup” phase, $R = P^H$ is the corresponding restriction and D_c the Galerkin projected coarse grid operator $D_c = P^HDP$.

As is discussed in [7, 31], this algebraic multigrid approach for D can be transferred to one for Q if the interpolation P preserves spin structure in the sense that on the coarse grid we can partition the degrees of freedom per grid point into two groups corresponding to different spins and that we have $\Gamma_5 P = P\Gamma_5^c$, where Γ_5^c is diagonal with values ± 1 , depending on the spin on the coarse grid. Putting $Q_c = \Gamma_5^c D_c$ we then have

$$I - PQ_c^{-1}P^H Q = I - PD_c^{-1}\Gamma_5^c P^H \Gamma_5 D = I - PD_c^{-1}P^H D, \quad (2.9)$$

showing that the coarse grid error propagator for D is identical to the coarse grid error propagator for Q if we take the same P . Note that the construction of P in [1, 19, 31] is indeed spin structure preserving, while this is not the case for the “little Dirac” construction found in [26].

As a matter of fact the SAP smoothing used in DD- α AMG is identical for D and Q , as well, which can be shown by the following argument. Mathematically, one step of SAP is a product of block projections, i.e., the error propagator is given by

$$E_{\text{SAP}} := \prod_{i=1}^b (I - \underbrace{I_{\mathcal{L}_i} Q_i^{-1} I_{\mathcal{L}_i}^H}_{:=M_{Q_i}} Q), \quad (2.10)$$

where b is the number of subdomains, \mathcal{L}_i is the i -th subdomain of the lattice \mathcal{L} , $I_{\mathcal{L}_i}$ the trivial injection from \mathcal{L}_i into \mathcal{L} , and $Q_i := I_{\mathcal{L}_i}^H Q I_{\mathcal{L}_i}$ the block restriction of Q on \mathcal{L}_i .

Note that algorithmically, the calculations corresponding to $I_{\mathcal{L}_i} Q_i^{-1} I_{\mathcal{L}_i}^H Q$ can be performed in parallel for all blocks i of the same color if we introduce a red-back ordering on the blocks.

With this we get the following proposition, in which we define M_{D_i} and D_i analogously to M_{Q_i} .

Proposition 1. *The error propagator $E_{\text{SAP}}(Q) := \prod_{i=1}^b (I - M_{Q_i} Q)$ is equivalent to $E_{\text{SAP}}(D) := \prod_{i=1}^b (I - M_{D_i} D)$.*

Proof. We first note that Γ_5 is just a local positive or negative identity, so its block restriction $\Gamma_5^i := I_{\mathcal{L}_i}^H \Gamma_5 I_{\mathcal{L}_i}$ on \mathcal{L}_i not only satisfies $I_{\mathcal{L}_i}^H \Gamma_5 = \Gamma_5^i I_{\mathcal{L}_i}^H$ but also

$(\Gamma_5^i)^{-1} = \Gamma_5^i$. To prove the proposition we only need to show that the error propagators are identical for any given subdomain i :

$$\begin{aligned}
I - M_{Q_i}Q &= I - (I_{\mathcal{L}_i}Q_i^{-1}I_{\mathcal{L}_i}^H)Q \\
&= I - (I_{\mathcal{L}_i}(I_{\mathcal{L}_i}^H\Gamma_5DI_{\mathcal{L}_i})^{-1}I_{\mathcal{L}_i}^H)\Gamma_5D \\
&= I - (I_{\mathcal{L}_i}(\Gamma_5^iD_i)^{-1}I_{\mathcal{L}_i}^H)\Gamma_5D \\
&= I - (I_{\mathcal{L}_i}D_i^{-1}\Gamma_5^iI_{\mathcal{L}_i}^H)\Gamma_5D \\
&= I - (I_{\mathcal{L}_i}D_i^{-1}I_{\mathcal{L}_i}^H\Gamma_5)\Gamma_5D = I - M_{D_i}D.
\end{aligned} \tag{2.11}$$

□

The proposition states that SAP for Q is equivalent to SAP for D if the block inversions for the block systems Q_i are performed exactly, which together with (2.9) implies that the DD- α AMG method has the same error propagator, irrespective of whether it is applied to Q or to D . As observed in [19] SAP smoothing works well for the standard Wilson-Dirac operator D thus it also works well for the Hermitian Wilson-Dirac operator Q . However, if we perform only approximate block inversions in SAP—and this is what one typically does—this situation becomes less clear; see Section 4.

Alternatively, instead of SAP one can use (restarted) GMRES as a smoother for Q . For the non-Hermitian Wilson-Dirac operator D this is used in the multigrid methods from [1, 6, 31], and since GMRES is also one of the most numerically stable Krylov subspace methods for indefinite systems, it is to be expected to work well as a smoother in a multigrid method for Q as well. Interestingly, for GMRES smoothing a connection between Q and D similar to what has just been exposed for SAP smoothing does not hold. We compare the above options for the smoothing method experimentally in Section 4.

3. Eigensolver

The *generalized Davidson* (GD) method [10, 28, 34] is an eigensolver framework which can be seen as a generalization of Arnoldi’s method. Its advantage is that it does not rely on a Krylov subspace structure and thus offers a more flexible way of steering the search space \mathcal{V}_m into a desired direction. The method successively generates a set of orthogonal vectors v_1, v_2, \dots, v_m , which span the search space \mathcal{V}_m . An approximate eigenpair (u, θ) with $u \in \mathcal{V}_m$ is chosen such that the Ritz-Galerkin condition

$$Au - \theta u \perp \mathcal{V}_m \tag{3.1}$$

holds, which amounts to solving the (small and dense) $m \times m$ eigenvalue problem

$$(V_m^H AV_m) s - \theta s = 0, \text{ where } V_m = [v_1 \mid \dots \mid v_m], \tag{3.2}$$

and then taking $u = V_m s$ with s the eigenvector from (3.2) whose eigenvalue θ is closest to the target eigenvalue. The search space is then extended by a

new vector t which is obtained as a function of the matrix A , the approximate eigenvalue θ , and the eigenvector residual $r := Au - \theta u$. The new vector v_{m+1} is then retrieved after orthogonalizing t against v_1, \dots, v_m and normalizing it.

For our work we focus on obtaining t as an (approximate) solution of the correction equation

$$(A - \tau I)t = r, \quad (3.3)$$

where τ is an estimate for the target eigenvalue. The choice of τ steers the expansion of the search space, and with it the Ritz values, towards the desired eigenvalue regions, e.g., eigenvalues with smallest absolute value or with largest imaginary part.

For the (Hermitian) Wilson-Dirac operator, Davidson-type methods are to be preferred over Arnoldi’s method due to the fact that Arnoldi’s method would require exact solves of the correction equation to maintain its constitutive orthogonality relations, whereas Davidson-type methods are tailored to accommodate approximate solutions, and these can be computed efficiently via some steps of multigrid preconditioned flexible GMRES.

A description of the generalized Davidson method to obtain one eigenpair is given in Algorithm 1. Techniques for computing several eigenpairs and restarting will be reviewed in the subsequent section.

Algorithm 1: Generalized Davidson (basic)

input: initial guess t , desired accuracy ε
output: eigenpair (λ, x)

- 1 $V = \emptyset$
- 2 **for** $m = 1, 2, \dots$
- 3 $t = (I - VV^H)t$
- 4 $v_m = t / \|t\|_2$
- 5 $V = [V \mid v_m]$
- 6 $H = V^H AV$
- 7 get target eigenpair (θ, s) of H
- 8 $u = Vs$
- 9 $r = Au - \theta u$
- 10 **if** $\|r\|_2 \leq \varepsilon$
- 11 $\lambda = \theta, x = u$
- 12 **return**
- 13 compute t as a function of A, r and θ

3.1. GD- λ AMG

The method we propose for the Hermitian Wilson-Dirac operator is based on Algorithm 1 but incorporates several adaptations for the Hermitian Wilson-Dirac operator and the underlying DD- α AMG multigrid solver.

A first challenge is that we are confronted with a “maximally indefinite” interior eigenvalue problem, seeking the eigenvalues closest to zero, while the

operator has a nearly equal amount of positive and negative eigenvalues. The basic generalized Davidson method uses the Rayleigh Ritz procedure to determine the Ritz approximation by solving the standard eigenvalue problem (3.2) for $H = V_m^H A V_m$. Ritz values approximate outer eigenvalues better and faster than the interior ones [34], which is why we use *harmonic Ritz values* [32] instead.

Definition 1 (Harmonic Ritz Values). *A value $\theta \in \mathbb{C}$ is called a harmonic Ritz value of A with respect to a linear subspace \mathcal{V} if θ^{-1} is a Ritz value of A^{-1} with respect to \mathcal{V} .*

As the exterior eigenvalues of A^{-1} are the inverses of the eigenvalues of A of small modulus, harmonic Ritz values tend to approximate small eigenvalues well. Inverting A to obtain harmonic Ritz values can be avoided with an appropriate choice for \mathcal{V} as stated in the following theorem; cf. [36].

Theorem 1. *Let \mathcal{V} be some m -dimensional subspace with basis v_1, \dots, v_m . A value $\theta \in \mathbb{C}$ is a harmonic Ritz value of A with respect to the subspace $\mathcal{W} := A\mathcal{V}$, if and only if*

$$A u_m - \theta u_m \perp A\mathcal{V} \text{ for some } u_m \in \mathcal{V}, u_m \neq 0. \quad (3.4)$$

With

$$V_m := [v_1 | \dots | v_m], W_m := A V_m \text{ and } H_m := (W_m^H V_m)^{-1} W_m^H A V_m,$$

(3.4) is equivalent to

$$H_m s = \theta s \text{ for some } s \in \mathbb{C}^m, s \neq 0 \text{ and } u_m = V_m s.$$

Due to Theorem 1, we can obtain harmonic Ritz values by solving the generalized eigenvalue problem

$$W_m^H A V_m u = \theta W_m^H V_m u. \quad (3.5)$$

The computational overhead compared to the standard Ritz procedure is dominated by m^2 additional inner products to build $W_m^H A V_m$. In our numerical tests, we have observed that this is compensated by a faster convergence of the generalized Davidson method, cf. Section 4.

Although the multigrid approach is viable for the Hermitian Wilson-Dirac operator Q , it is, in practice, slower than for D . For exact solves of the subdomain systems in the SAP smoother, the discussion in Section 2 and Proposition 1 implies that the convergence speeds for Q and for D are comparable as the error propagation operators are identical. Though in computational practice, it is more efficient to do only approximate solves for the subdomain systems, using a small number of GMRES steps, for example. In this scenario the multigrid method becomes significantly slower when used for Q rather than D , see Figure 6 in Chapter 4. This slowdown can be countered by left-preconditioning the

correction equation with Γ_5 . This means that instead of solving (1.1) with Q , we can transform it equivalently according to

$$(Q - \tau I)t = r \tag{3.6}$$

$$\iff \Gamma_5(Q - \tau I)t = \Gamma_5 r$$

$$\iff (D - \tau \Gamma_5)t = \Gamma_5 r. \tag{3.7}$$

The spectrum of the resulting operator $\Gamma_5 Q(\tau) := D - \tau \Gamma_5$ has similarities to that of D with some eigenvalues collapsing on the real axis. As we will see in Chapter 4, this simple transformation speeds up the multigrid method significantly. Figure 2 shows full spectra of D , Q and $\Gamma_5 Q(\tau)$ for a configuration on a small 4^4 lattice.




Figure 2: Full Spectra of D , Q and $\Gamma_5 Q(\tau)$ for configuration 5 (see Table 1). The shaded area of $\text{spec}(Q)$ highlights the eigenvalues we are particularly interested in.

Restarting and locking. As the search space grows in every outer iteration, the storage and orthogonalization costs of the outer iteration in a generalized Davidson method eventually become prohibitively large. The following techniques reduce these costs in order to achieve a near-linear scaling in the number of computed eigenpairs. The first technique is *thick restarting* [36]. When the search space reaches a size of m_{max} , we perform a restart by discarding the current search space. At the same time we keep the first m_{min} smallest non-converged harmonic Ritz vectors and use them to span the search space at the beginning of the next restart cycle. We tuned the parameters m_{min} and m_{max} such that we have reason to assume that we keep both positive and negative harmonic Ritz values within the new search space. This way the eigensolver obtains a (nearly) equal amount of positive and negative eigenpairs in a uniform way.

In order to avoid re-targeting converged eigenpairs, we employ the concept of *locking* converged eigenpairs [40] as a second technique. Locking keeps the search space \mathcal{V} orthogonal to the space of already converged eigenvectors \mathcal{X} . In this manner, it is not required to keep converged eigenvectors in the search space which has the effect that the search space dimension becomes bounded independently of the number of eigenpairs sought. This in turn bounds the cost for computing the harmonic Ritz pairs. The new search direction still has to be orthogonalized against all previous eigenvectors, which leads to costs of order

$\mathcal{O}(nk^2)$, since it consists of $k - 1$ inner products for each of the k eigenpairs. This is responsible for the fact that, in principle, the cost of our method scales superlinearly with k , and this becomes visible when k becomes sufficiently large.

Local coherence and its effect on the correction equation. The strength of algebraic multigrid methods relies on an effective coarse grid correction step and thus on the construction of the interpolation operator P . The methods in use for the Wilson-Dirac operator are all adaptive: They require a setup phase which computes “test vectors” $w_i, i = 1, \dots, n_{tw}$ which are collected as columns in the matrix $W = [w_1 | \dots | w_{n_{tw}}]$. The test vectors are approximations to eigenvectors corresponding to small eigenvalues of the unshifted Wilson-Dirac operator D . The matrix W is then used to build an aggregation based, “block diagonal” interpolation operator P , where each diagonal block constitutes an *aggregate*, i.e., a block \mathcal{A}_i of W corresponding to the degrees of freedom of a block of the lattice \mathcal{L} ; see Figure 3 and [19, 31].




Figure 3: Matrix view of the construction of the aggregation based interpolation operator P .

By construction, the range of an aggregation based interpolation P contains *at least* the range spanned by the test vectors it is being built from. In [26] it has been observed that eigenvectors belonging to small eigenvalues of the Wilson-Dirac operator D are *locally coherent* in the sense that these eigenvectors are locally similar, i.e., they are similar on the individual aggregates. This is the reason why the span of an aggregation based interpolation P contains good approximations to small eigenpairs *far beyond* those which are explicitly used for its construction. This in turn explains the efficiency of such P in the multigrid method.

We can study local coherence using the *local coherence measure* lc of a vector v defined as

$$lc(v) = \|\Pi v\|/\|v\|,$$

where Π denotes the orthogonal projection on the range of P . If $lc(v)$ is close to 1, there is a good approximation to v in the range of P , implying that the multigrid coarse grid correction reduces error components in the direction of v almost to zero.

Figure 4 gives values for $lc(v)$ for the Wilson-Dirac operator D and the corresponding Hermitian Wilson-Dirac operator Q on a 4^4 lattice. Since this lattice is so small, we can compute the full spectrum ($12 \cdot 4^4 = 3072$ eigenpairs) of both D and Q . For each matrix we then consider a partitioning of the




Figure 4: Local coherence for D (left) and Q (right) for a 4^4 configuration, cf. Table 1.

eigenvectors into 128 sets, each set consisting of 24 consecutive eigenpairs. Here, “consecutive” refers to an ordering based on the modulus and the sign of the real part; see the next paragraph for details. For each of these “interpolation sets”, the corresponding row displays the color coded value of $lc(v)$ when projecting an eigenvector v with the projection Π corresponding to the aggregation-based interpolation P built with the eigenvectors from that interpolation set as test vectors. The aggregates used were based on a decomposition of the 4^4 lattice into 16 sub-lattices of size 2^4 . Due to the spin structure preserving approach, we have two aggregates per sub-lattice, each built from the corresponding spin components of the 24 test vectors³. Of course $lc(v) = 1$ (dark red) if v is from the respective interpolation set.

The numbering of the eigenvalues used in these plots is as follows: The plot for D has its eigenvalues with negative imaginary part in its left half, ordered by descending modulus and enumerated by increasing, negative indices including zero, $-1\ 535, \dots, 0$. The eigenvalues with positive imaginary part are located in the right half, ordered by ascending modulus and enumerated with increasing positive indices $1, \dots, 1536$. For Q we just order the real eigenvalues by the natural ordering on the reals, using again negative and positive indices. Thus, for D as for Q , eigenvalues small in modulus are in the center and their indices are small in modulus, while eigenvalues with large modulus appear at the left and right ends and their indices are large in modulus.

Although one must be careful when drawing conclusions from extremely small configurations, Figure 4 illustrates two important phenomena. Firstly, local coherence appears for both D and Q , but it is more pronounced for the non-Hermitian Wilson-Dirac matrix. This especially holds directly next to the interpolation sets (the diagonal in the plots). Secondly, local coherence is par-

³The projection Π therefore projects onto a subspace of dimension $24 \cdot 2 \cdot 16 = 768$. If there were no local coherence at all, the expected value of lc is thus $768/(12 \cdot 4^4) = 0.25$.

ticularly strong and far-reaching when projecting on the interpolation sets corresponding to the smallest and largest eigenpairs in absolute values. In the center of both plots, we observe a star-shaped area with particularly high local coherence. This area corresponds to around 10% of the smallest eigenvalues. To a lesser extent, local coherence is also noticeable for the other parts of the spectrum, as we consistently observe higher values for $lc(v)$ for eigenvectors close to the respective interpolation set.

The right part of Figure 5 reports similar information for the Hermitian Wilson-Dirac operator Q coming from a larger, realistic configuration on a 64×32^3 lattice. For lattices of this size we cannot compute the full spectrum, thus we show the values for the 984 smallest eigenpairs, subdivided in 41 interpolation sets, each consisting of 24 consecutive eigenpairs. The aggregates were this time obtained from 4^4 sublattices. For comparison, the left part of the figure shows a zoomed-in part of the local coherence plot for Q for the 4^4 -lattice from Figure 4.




Figure 5: Local coherence for Q for different lattices, focusing on the eigenpairs closest to zero. *Left*: 432 eigenpairs of a 4^4 lattice. *Right*: 984 eigenpairs of a 64×32^3 lattice.

First note that the colors encode different values in the left and right part of Figure 5. Local coherence does not drop below 0.9 for the large configuration, while for the small configuration it goes down to 0.6. On the other hand, 984 eigenvalues only correspond to a minuscule fraction of roughly $4 \cdot 10^{-3}\%$ of the total of $12 \cdot 32^3 \cdot 64 = 25\,165\,824$ eigenvalues, which is much less than the roughly 10% depicted for the small configuration. For the interpolation operator of the large configuration we used aggregates corresponding to 4^4 sub-lattices, which give a total of $2 \times 2^{13} = 16\,384$ aggregates. In relative terms, this is several orders of magnitude finer as for the 4^4 lattice. This finer aggregation leads to interpolation operators with increased faculties to recombine information, which explains the resulting higher local coherence.

Both parts of Figure 5 show that local coherence drops off for eigenvectors further away from the interpolation set. For the large configuration, we see, for example, that local coherence of the vectors from the last interpolation set with the second-to-last interpolation set is very high as indicated by the deep red color in the top right corner of the plot. The local coherence of these vectors

with respect to the central interpolation set (which contains the eigenpairs with eigenvalues closest to 0) is significantly smaller, indicated by the yellow color at the middle of the right-hand boundary of the plot. In the scenario where we choose the shift τ in the correction equation (3.7) farther away from zero—as we are targeting eigenpairs close to τ —a coarse grid operator constructed using the eigenpairs closest to zero thus becomes increasingly less effective, reducing the overall convergence speed of the multigrid method significantly. To remedy this, we propose a dynamical interpolation updating approach, resulting in a coarse grid operator that remains effective on the span of the eigenvectors with eigenvalues close to the value of τ set in the outer iteration of the generalized Davidson method. In the course of the outer iteration, once enough eigenpairs are available, we therefore rebuild the interpolation, and with it the coarse grid operator, using the already converged eigenvectors which are *closest* to the currently targeted harmonic Ritz value. Once a harmonic Ritz value converged to an eigenvalue and we choose a new target value τ that has the same sign as the previous target, we replace one eigenvector from the interpolation set—the one farthest away from τ —with the newly converged eigenvector, and update the multigrid hierarchy. If the new τ has its sign opposite to the previous one we replace the full interpolation set with converged eigenvectors closest to the new τ , and again update the multigrid hierarchy. The updates of the interpolation and coarse grid operators involve some data movement and local operations, but their cost is minor compared to the cost of the other parts of the computation.

With this approach, the coarse grid is always able to treat the eigenspace closest to our current harmonic Ritz approximation efficiently and makes optimal use of the existing local coherence. This results in a significantly faster multigrid solver when larger shifts are used, i.e., when a large number of small eigenpairs has to be computed. Since the solution of these shifted systems accounts for most of the work in the eigensolver this approach improves the eigenvalue scaling to a nearly linear complexity, as seen in Section 4.

Algorithm 2: GD- λ AMG

input: Hermitian Dirac operator Q , no. of eigenvalues n , no. of test vectors n_{tv} , min. and max. subspace size m_{min} and m_{max} , initial guess $[v_1, \dots, v_{n_{tv}}, t] =: [V|t]$, desired accuracy ε_{outer}

output: set of n eigenpairs (Λ, X)

```

1  $\Lambda = \emptyset, X = \emptyset$ 
2 for  $m = n_{tv} + 1, n_{tv} + 2, \dots$ 
3    $t = (I - VV^H)t, t = (I - XX^H)t$ 
4    $v_m = t/\|t\|_2$ 
5    $V = [V|v_m]$ 
6   get all  $(\theta_i, s_i)$  with  $(QV)^H(QV)s_i = \theta_i(QV)^HVs_i$ 
7   find smallest (in modulus)  $\theta_i \notin \Lambda$ 
8    $u = Vs_i, r = Qu - \theta_i u$ 
9   if  $\|r\|_2 \leq \varepsilon_{outer}$ 
10     // current eigenpair has converged
11      $\Lambda = [\Lambda, \theta_i], X = [X, u]$ 
12     update smallest (in modulus)  $\theta_i \notin \Lambda$ 
13      $u = Vs_i, r = Qu - \theta_i u$ 
14     rebuild interpolation using the  $n_{tv}$  eigenvectors  $x_j$  with
15     eigenvalue  $\lambda_j$  closest to  $\theta_i$  and update multigrid hierarchy
16     // solve correction equation
17      $t = \text{DD-}\alpha\text{AMG}((D - \theta_i \Gamma_5), \Gamma_5 r)$ 
18     // restart
19     if  $m \geq m_{max}$ 
20       get  $(\Theta, S)$  as all eigenpairs  $(\theta_i, s_i)$  of
21        $(QV)^H(QV)s_i = \theta(QV)^HVs_i$ 
22       sort  $(\Theta, S)$  by ascending modulus of  $\Theta$ 
23       for  $i = 1, \dots, m_{min}$ 
24          $V_i = Vs_i$ 
25          $(QV)_i = (QV)s_i$ 
26       retain first  $m_{min}$  vectors of  $V$  and  $QV$ 

```

A summary of our eigensolver, termed GD- λ AMG, a generalized Davidson method with algebraic multigrid acceleration, is given as Algorithm 2.

4. Numerical Tests

In this section we present a variety of numerical tests to analyze the efficiency of the GD- λ AMG eigensolver.

Table 1 contains information on the gauge configurations we use in our tests. The two small configurations on a 4^4 and a 8^4 lattice were generated using our own heat-bath algorithm. The configurations on the larger lattices (configurations 1 to 4) were provided by our partners at the University of Regensburg within the Collaborative Research Centre SFB-TRR55; see [3]. For these con-

| ID | lattice size $N_t \times N_s^3$ | hopping parameter κ | mass parameter $a \cdot m_0$ | clover term c_{sw} | CPU cores |
|----|------------------------------------|-------------------------------|---------------------------------|-------------------------|--------------|
| 1 | 48×24^3 | 0.13620 | -0.3289 | 1.9192 | 648 |
| 2 | 64×32^3 | 0.13632 | -0.3322 | 1.9192 | 1,024 |
| 3 | 64×40^3 | 0.13632 | -0.3322 | 1.9192 | 2,000 |
| 4 | 64×64^3 | 0.13632 | -0.3322 | 1.9192 | 4,096 |
| 5 | 4×4^3 | – | -0.7867 | 0 | – |
| 6 | 8×8^3 | – | -0.7972 | 0 | – |

Table 1: Configurations used for numerical tests. Configurations 1–4 are configurations from the ensembles II and IV–VI from [3]. Configurations 5 and 6 were generated locally and are mainly used for small scale MATLAB experiments.

| | parameter | symbol | default |
|------------------------|-------------------------------|-----------------------|-------------------|
| DD- α AMG setup | number of test vectors | n_{tv} | 24 |
| | setup iterations | | 6 |
| | (post-)smoothing steps | | 4 |
| DD- α AMG solve | relative residual | ε_{inner} | 10^{-1} |
| | maximum iterations | | 5 |
| | coarse grid tolerance | | $5 \cdot 10^{-1}$ |
| eigensolver method | relative eigenvector residual | ε_{outer} | 10^{-8} |
| | number of eigenpairs | | 100 |
| | minimum subspace size | m_{min} | 30 |
| | maximum subspace size | m_{max} | 50 |

Table 2: List of algorithmic *default* parameters.

figurations, we actually use *clover improved* (see [35]) Wilson-Dirac operators D and Q , where a block diagonal term with 6×6 diagonal blocks is added to improve the lattice discretization error from $\mathcal{O}(a)$ to $\mathcal{O}(a^2)$. The resulting modified D is still Γ_5 -Hermitian. The mass parameter m_0 from (2.5) is chosen such that $a \cdot m_0 = \frac{1}{2\kappa} - 4$ for computations with configurations 1 to 4. For the small configurations we chose m_0 such that we obtain a comparable conditioning of the matrix. A second table, Table 2, shows the default algorithmic parameter settings we used within GD- λ AMG. For a more detailed explanation of these parameters we refer to [19].

The numerical results involving configurations 1–4 were obtained on the JU-RECA and JUWELS clusters at the Jülich Supercomputing Centre [23, 24], while results involving the other lattices were obtained on a smaller workstation. We will compare our results with the state-of-the-art library PRIMME and with PARPACK, and we start by outlining their underlying basic algorithms.

PRIMME (PREconditioned Iterative MultiMethod Eigensolver) [41, 43] implements a broad framework for different Davidson-type eigensolvers. Its performance is best if it is given an efficient routine to solve linear systems with the matrix A , and we do so by providing the Γ_5 -preconditioned DD- α AMG solver. There are two key differences compared to GD- λ AMG:

- The interpolation cannot be updated efficiently within the PRIMME framework (at least not without expert knowledge on the underlying data structures), hence we do not update it for this method.
- PRIMME uses a Rayleigh Ritz instead of a harmonic Ritz approach to extract eigenvalue approximations.

PRIMME has a fairly fine-tuned default parameter set, e.g., for subspace size or restart values, and is able to dynamically change the eigensolver method. We keep the default settings and provide the same multigrid solver to PRIMME as we do for GD- λ AMG.

PARPACK (Parallel ARnoldi PACKage) [38] is a somewhat older but widely used software for the computation of eigenvalues of large sparse matrices. It is based on an implicitly restarted Arnoldi method, which is originally designed to find extremal eigenvalues. It is possible to transform an interior problem into an exterior one using a filter polynomial, i.e., a polynomial which is large on the k interior eigenvalues we are looking for and small on the remaining ones. To construct such a polynomial, for example as a Chebyshev polynomial, we need information on the eigenvalue λ_{max} which is largest in modulus and the $(k + 1)$ st smallest in modulus, λ_{k+1} . While $\lambda_{max} = 8$ is a sufficiently good estimate for the Hermitian Wilson-Dirac matrix Q , no a-priori guess for λ_{k+1} is available in realistic scenarios. For our tests, we run one of the other methods to compute the first k eigenvalues and then use a slightly larger value as a guess for λ_{k+1} . While this approach obviously costs a lot of additional work and actually makes the subsequent Arnoldi method obsolete, it is a good reference for a *near-optimally* polynomially filtered Arnoldi method. Since this approach does not require inversions of the matrix Q , the parameter set for this method is rather small. We use a degree ten Chebyshev polynomial as the filter polynomial and set the maximum subspace size to be twice the number of sought eigenpairs. The required eigenvector residual is set to 10^{-8} , as with the other methods.

4.1. Algorithmic tuning

Solving the correction equation.. Each step of Algorithm 2 uses DD- α AMG in line 14 to solve the Γ_5 -preconditioned correction equation $(D - \tau\Gamma_5)t = \Gamma_5r$. More precisely, as indicated by the parameters given in the middle of Table 2, we stop the outer (FGMRES) iteration of DD- α AMG once the initial residual norm is reduced by a factor of 0.1 or a maximum of 5 iterations is achieved. Within each DD- α AMG iteration we require a reduction of the residual by a

| correction equation | iterations | | Time in core-h. |
|---|------------|--------|--------------------|
| | outer | inner | |
| Eq. (3.6): $(Q - \tau I)t = r$ | 565 | 10,349 | 83.0 |
| Eq. (3.7): $(D - \tau\Gamma_5)t = \Gamma_5 r$ | 511 | 3,045 | 41.3 |

Table 3: Impact of Γ_5 -preconditioning for the computation of 100 eigenpairs of configuration 1 (see Table 1).

factor of 0.5 when solving the system on the coarsest level. Table 3 shows that the Γ_5 -preconditioning yields indeed significant gains in compute time.

A variant of generalized Davidson methods solves, instead of the correction equation (3.3), the Jacobi-Davidson projected [36] system $(I - uu^H)(A - \theta I)(I - uu^H)$, where u is the last (harmonic) Ritz vector approximation. This will avoid stagnation in the case that the correction equation is solved *too exactly*. There are theoretically justified approaches which adaptively determine how accurately the projected system should be solved in each iteration. Since we solve the correction equation to quite low relative precision (10^{-1} only), we could not see a benefit from using the Jacobi-Davidson projected system. Indeed, even with the adaptive stopping criterion, this approach increased the compute time by approximately 15%.

Impact of the smoother. The original DD- α AMG method uses SAP as a smoother and we have shown in Chapter 2 that SAP is also applicable for the Hermitian Wilson-Dirac operator Q , yielding the same error propagation operator as long as the individual block systems are solved exactly. We now compare a cost-efficient, approximate SAP and GMRES as smoothers within the multigrid methods constructed for the matrices $D - \tau I$, $Q - \tau I$ and $D - \tau\Gamma_5$, where τ ranges from 0 to 0.5 for configuration 6. Note that $D - \tau I$ is not relevant for this work, since it would arise when computing eigenpairs for D . We still include the results here to be able to compare the performance of DD- α AMG for $Q - \tau I$ and $D - \tau\Gamma_5$ with the performance of DD- α AMG for $D - \tau I$.

Figure 6 shows a scaling plot with respect to the target shift τ for configuration 6. For this plot, we used a two-level DD- α AMG method with six steps of the adaptive setup procedure to generate the coarse grid system. The GMRES smoother requires a reduction of the residual by a factor of 10^{-1} or after a maximum of ten iterations have been performed. Similarly, the SAP smoother performs three sweeps of SAP, where each block solve is performed using GMRES until a reduction of the residual by a factor of 10^{-1} was achieved for the individual block or after a maximum of ten iterations have been performed. This way the computational work for both smoothers is roughly comparable.

Figure 6 verifies what was stated in Section 3.1, namely that DD- α AMG converges more slowly for Q compared to D . It also shows that Γ_5 -preconditioning is beneficial in the case of GMRES smoothing whereas in the case of SAP smoothing, it loses efficiency compared to Q , although only by a small margin.




Figure 6: Comparison of iteration counts of the DD- α AMG method using either SAP or GMRES smoothing for configuration 6 and increasing target shifts τ . The black diamonds at the bottom depict the eigenvalue distribution of Q .

Comparing the two smoothing methods for $D - \tau\Gamma_5$, we see that both methods perform nearly identical up to larger shifts, where SAP starts to be slightly more favorable. We do not expect this to be relevant for larger configurations, though, since there the spectrum is much more dense. Even when aiming for a large number of small eigenvalues, we certainly do not expect to end up with τ -values as large as 0.2 already. Since the focus of this paper is on finding an efficient coarse grid operator, and not on optimizing the smoother, we stick to GMRES smoothing here. Implementing SAP instead of GMRES for $D - \tau\Gamma_5$ within the DD- α AMG framework would require a more substantial remodeling of the DD- α AMG code.

Impact of the coarse grid correction. For an assessment of the impact of the coarse grid correction step we compute 100 eigenvalues for configuration 1, once using DD- α AMG with GMRES smoothing to solve the correction equation, and once with a modification where we turned-off the coarse grid correction. This yields a generalized Davidson method where the Γ_5 -preconditioned correction equation (3.7) is solved using FGMRES with the GMRES-steps of the smoother as a non-stationary preconditioner, i.e., GMRESR, the recursive GMRES method [42]. Note that we do not yet include updating the multigrid hierarchy as the outer iterations proceeds.

The left part of Figure 7 shows the FGMRES iterations spent on the correction equation for computing 100 eigenvalues for the two variants. We see that right from the beginning, including the coarse grid correction, i.e., using the multigrid method, reduces the iteration count by one order of magnitude com-

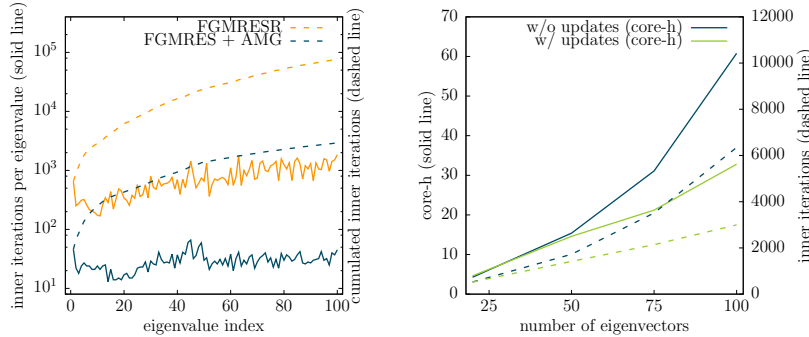


Figure 7: *Left*: Computation of 100 eigenvalues for configuration 1 with GMRESR and FGMRES + AMG. *Right*: Comparing eigenvalue scaling for configuration 2 depending on whether eigenvalue information is provided for the interpolation operator.

pared to the “pure” GMRESR-Krylov subspace method. The required number of FGMRES iterations per eigenvalue stays constant at ≈ 30 for the multigrid method, whereas GMRESR starts at ≈ 300 and increases to $\approx 1,200$ for the last eigenvalues. This is also reflected in CPU time, where on JUWELS multigrid preconditioning results in 30 core-h for the entire computation, whereas 217 core-h were necessary when using GMRESR. Thus multigrid gains one order of magnitude and, in addition, shows an improved scaling behaviour, despite the loss of local coherence for the larger eigenvalues.

The right part of Figure 7 now illustrates the additional benefits that we get from turning on the updating of the multigrid hierarchy, i.e., when performing full GD- λ AMG as described in Algorithm 2. Both approaches perform similarly as long as a small amount of eigenvalues is sought. This changes substantially for already a moderate amount of eigenvalues to a point where interpolation updates save roughly a factor of two in both, number of iterations (dashed lines) and consumed core-h (solid lines). In terms of iterations it is also noteworthy that interpolation updates lead to a nearly linear scaling with respect to the eigenvalue count, whereas in the other case the scaling is closer to quadratic.

4.2. Scaling results

Scaling with the lattice size.. We now compare GD- λ AMG, PRIMME and PARPACK in terms of scaling with respect to the lattice size. For this, we report the total core-h consumed for computing 100 eigenpairs on configurations 1 to 4.

Figure 8 shows that PRIMME and GD- λ AMG scale similarly with increasing lattice size. GD- λ AMG shows some improvement in core-h compared to PRIMME, and this improvement tends to get larger when increasing the lattice size. The right part of Figure 8 shows, that this improvement might be partially attributed to the fact that we use a harmonic Ritz extraction. Here, we compare GD- λ AMG with its default harmonic Ritz extraction to a variant where we use the standard Rayleigh-Ritz extraction as is done in PRIMME. The Figure shows that harmonic Ritz extractions result in substantially less inner




Figure 8: *Left*: Computation of 100 eigenvalues for 48×24^3 to 64×64^3 lattices for different methods. *Right*: Comparison of Ritz and harmonic Ritz eigenpair extraction for different lattice sizes.

iterations. This also yields savings in computational time, which are smaller, due to the additional cost for the inner products. Note that for larger lattices eigenvalues become more clustered. The harmonic Ritz extraction is then more favorable compared to the Rayleigh-Ritz approach, since it is able to better separate the target eigenvalue from the neighboring ones. PARPACK scales worse than the other methods, even when we use an unrealistic “near optimal” filter polynomial as we did here. In practice, i.e., when no guess for $|\lambda_{k+1}|$ is available, PARPACK’s performance would fall even further behind. Applying PARPACK to Q^{-1} to make use of the efficient multigrid solver is way too costly, due to the necessity of very exact solves to maintain the Krylov structure.

Scaling with the number of eigenvalues. Figure 9 reports results of a scaling study obtained for configuration 2. We just compare GD-λAMG and PRIMME, since PARPACK is not competitive.

The figure shows that GD-λAMG has an advantage over PRIMME when larger numbers of eigenvalues are sought. GD-λAMG needs up to one order of magnitude less iterations, which translates to a speed-up of 1.5 for 50 eigenvalues to up to more than three for 1000 eigenvalues. This shows that the additional effort due to the adaptive construction of the multigrid hierarchy and the harmonic Ritz extraction is beneficial with respect to the overall performance. GD-λAMG and PRIMME both scale nearly linearly with respect to the number of eigenvalues sought, up to at least 300 eigenvalues. Then PRIMME’s performance starts to decrease more significantly compared to GD-λAMG. We see that the increase in the overall computing time in PRIMME scales more than linearly with the number of iterations to be performed. This indicates that the non-adaptive multigrid solver used in PRIMME is getting increasingly less efficient, a situation that is remedied with the update strategy realized in GD-λAMG.




Figure 9: Eigenvalue scaling in the range of 50 to 1000 eigenpairs on configuration 2 with a lattice size of 64×32^3 .

5. Spectral gap

This paper has a clear focus on algorithmic development to compute small eigenmodes. For purposes of illustration, we now include an example which relates our computed eigenvalues to the general approach of lattice QCD.

The *spectral gap* is the relevant quantity for the stability of Monte Carlo simulations of lattice QCD [13]. It is defined as the smallest eigenvalue in magnitude of the Hermitian Wilson-Dirac operator Q . If chiral symmetry was preserved, the spectral gap would be bounded from below by the bare current-quark mass m . Since the Wilson-Dirac operator breaks chiral symmetry it is possible that the gap is smaller than m . In the following we adopt the notation of [13] and denote the eigenvalues of Q^2 by α_i .

Figure 10 shows the number of eigenvalues $\sqrt{\alpha_i}$ per bin and fm^4 (blue line) at the low end of the spectrum computed on configuration 2 with a lattice size of 64×32^3 and $N_f = 2$ mass-degenerate non-perturbatively $\mathcal{O}(a)$ improved quarks, cf. Table 1. The lattice spacing is $a = 0.071 \text{ fm}$ from [3] and the bin size is set to 1 MeV.

It is interesting to compare these computational results to what we know analytically in the continuum theory in the infinite volume limit. The spectral density can be computed from the Banks–Casher relation [4, 13] asymptotically

$$\tilde{\rho}(\sqrt{\alpha}) \stackrel{\alpha > m^2}{\approx} 2\sqrt{\alpha} \left[\frac{\Sigma}{\pi\sqrt{\alpha - m^2}} + \mathcal{O}(1) \right], \quad (5.1)$$

which goes to infinity for $\sqrt{\alpha} \rightarrow m$ and is not defined for $\sqrt{\alpha} < m$, resulting in no eigenvalues smaller than m for the continuum case. Note that the Banks–Casher relation can be used to determine the chiral condensate Σ from the non-zero density of eigenmodes at the origin in infinite volume; cf. [22]. In




Figure 10: Number of eigenvalues $\sqrt{\alpha}_i$ of $|Q|$ per bin and fm^4 at the low end of the spectrum on configuration 2 with a lattice size of 64×32^3 . For comparison the same quantity in the continuum theory is also shown, cf. Eq. (5.1), taking $m = 8.22 \text{ MeV}$ and $\Sigma = (250 \text{ MeV})^3$.

order to compare to our lattice results for configuration 2 we compute the bare current-quark mass m using [17, Eq. (E.1)]. The spectral gap (black vertical line in figure 10) is given by $\sqrt{\alpha} = Z_A m$ [13], which we evaluate using the axial-current renormalization constant Z_A from [9]. The resulting distribution of eigenvalues $\sqrt{\alpha}$ per bin and fm^4 for the continuum operator (red line) is also plotted in Figure 10. The comparison to the lattice data shows that on configuration 2 there is a significant amount of eigenvalues smaller than $\sqrt{\alpha}$, and the distribution close to the gap deviates from Equation (5.1). The figure thus illustrates quantitatively the deviations due to lattice artifacts which are not unexpected for the given lattice sizes.

6. Conclusion

In this paper we introduced an eigensolver built around the multigrid method DD- α AMG for efficient shifted inversions within a generalized Davidson method. Several adaptations were included in order to improve eigenvalue scaling by at least a factor of three for a moderate to large amount of eigenvalues compared to current general purpose eigensolver software. This was accomplished by implementing a synergy between the generalized Davidson method and the AMG solver. Additionally, we incorporated several state-of-the-art techniques like locking, thick restarting and harmonic Ritz extraction to achieve a nearly linear eigenvalue scaling. We included a variety of numerical tests to verify the efficiency of our proposed adaptations.

We can now use the new algorithm to compute relatively many small eigenmodes of large configurations. This allows to advance deflation approaches

for stochastic trace estimation as required in the computation of disconnected fermion loops. We plan to explore this further in cooperation with the European twisted mass collaboration. Note that the eigenpairs of the symmetrized Wilson and twisted mass operators differ by an imaginary shift in the eigenvalues only.

With the algorithm presented in this work it is now also possible to perform a systematic study of the spectral gap and the spectral density varying the lattice volume, the lattice spacing, and the quark mass. We plan to extend the work of [13] to the $N_f = 2 + 1$ $O(a)$ improved theory, cf. [8] in the future.

6.1. Acknowledgment

The authors are grateful for the fruitful collaboration with Benjamin Müller of the University of Mainz and our colleagues within the SFB-TR 55 from the Theoretical Physics department of the University of Regensburg, especially Gunnar Bali and his group for providing configurations. We thank the Jülich Supercomputing Centre (JSC) for access to high performance computing resources through NIC grant HWU29. This work is mainly funded by the Deutsche Forschungsgemeinschaft (DFG) via the collaborative research center SFB-TR55.

References

- [1] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. A. Manteuffel, S. F. McCormick, J. C. Osborn, and C. Rebbi. Adaptive multigrid algorithm for the lattice Wilson-Dirac operator. *Phys. Rev. Lett.*, 105:201602, 2010.
- [2] G. Bali, S. Collins, A. Frommer, K. Kahl, I. Kanamori, B. Müller, M. Rottmann, and J. Simeth. (Approximate) Low-Mode Averaging with a new Multigrid Eigensolver. *PoS, LATTICE2015:350*, 2015.
- [3] G. Bali, S. Collins, B. Gläfle, M. Göckeler, J. Najjar, R. H. Rödl, A. Schäfer, R. W. Schiel, A. Sternbeck, and W. Söldner. The Moment $\langle x \rangle_{u-d}$ of the Nucleon From $N_f = 2$ Lattice QCD Down to Nearly Physical Quark Masses. *Phys. Rev.*, D90(7):074510, 2014.
- [4] T. Banks and A. Casher. Chiral Symmetry Breaking in Confining Theories. *Nucl. Phys. B*, 169:103–125, 1980.
- [5] B. Blossier, M. Della Morte, N. Garron, G. von Hippel, T. Mendes, H. Simma, and R. Sommer. HQET at Order $1/m$: II. Spectroscopy in the Quenched Approximation. *JHEP*, 05:074, 2010.
- [6] J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn, and C. Rebbi. Adaptive multigrid algorithm for lattice QCD. *Phys. Rev. Lett.*, 100:041601, 2007.
- [7] James Brannick and Karsten Kahl. Bootstrap algebraic multigrid for the 2d Wilson Dirac system. *SIAM Journal on Scientific Computing*, 36(3):B321–B347, 2014.

- [8] M. Bruno, D. Djukanovic, G. P. Engel, A. Francis, G. Herdoiza, H. Horch, P. Korcyl, T. Korzec, M. Papinutto, S. Schaefer, E. E. Scholz, J. Simeth, H. Simma, and W. Soeldner. Simulation of QCD With $N_f = 2 + 1$ Flavors of Non-Perturbatively Improved Wilson Fermions. *JHEP*, 02:043, 2015.
- [9] M. Dalla Brida, T. Korzec, S. Sint, and P. Vilaseca. High Precision Renormalization of the Flavour Non-Singlet Noether Currents in Lattice QCD With Wilson Quarks. *Eur. Phys. J. C*, 79(1):23, 2019.
- [10] E. R. Davidson. The Iterative Calculation of a few of the Lowest Eigenvalues and Corresponding Eigenvectors of Large Real-Symmetric Matrices. *Journal of Computational Physics*, 17:87–94, January 1975.
- [11] T. DeGrand and C. E. Detar. *Lattice Methods for Quantum Chromodynamics*. World Scientific, 2006.
- [12] T. A. DeGrand and S. Schaefer. Improving Meson two Point Functions in Lattice QCD. *Comput. Phys. Commun.*, 159:185–191, 2004.
- [13] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio, and N. Tantalo. Stability of Lattice QCD Simulations and the Thermodynamic Limit. *JHEP*, 02:011, 2006.
- [14] S. C. Eisenstat, H. C. Elman, and M. H. Schultz. Variational iterative methods for nonsymmetric systems of linear equations. *SIAM Journal on Numerical Analysis*, 20(2):345–357, 1983.
- [15] E. Endress, C. Pena, and K. Sivalingam. Variance Reduction With Practical all-to-all Lattice Propagators. *Comput. Phys. Commun.*, 195:35–48, 2015.
- [16] J. Foley, K. Jimmy Juge, A. O’Cais, M. Peardon, S. M. Ryan, and J. Skullerud. Practical all-to-all Propagators for Lattice QCD. *Comput. Phys. Commun.*, 172:145–162, 2005.
- [17] P. Fritzsche, F. Knechtli, B. Leder, M. Marinković, S. Schaefer, R. Sommer, and F. Virota. The Strange Quark Mass and Lambda Parameter of two Flavor QCD. *Nucl. Phys. B*, 865:397–429, 2012.
- [18] A. Frommer, K. Kahl, S. Krieg, B. Leder, and M. Rottmann. Aggregation-based multilevel methods for lattice QCD. In *Proceedings of Science*, <http://pos.sissa.it>, volume LATTICE2011:046, 2011.
- [19] A. Frommer, K. Kahl, S. Krieg, B. Leder, and M. Rottmann. Adaptive aggregation based domain decomposition multigrid for the lattice Wilson-Dirac operator. *SIAM J. Sci. Comp.*, 36(4):A1581–A1608, 2014.
- [20] C. Gattringer and C. B. Lang. *Quantum Chromodynamics on the Lattice*, volume 788 of *Lect. Notes Phys.* Springer, 2009.

- [21] L. Giusti, P. Hernandez, M. Laine, P. Weisz, and H. Wittig. Low-Energy Couplings of QCD From Current Correlators Near the Chiral Limit. *JHEP*, 04:013, 2004.
- [22] L. Giusti and M. Luscher. Chiral Symmetry Breaking and the Banks-Casher Relation in Lattice QCD With Wilson Quarks. *JHEP*, 03:013, 2009.
- [23] Jülich Supercomputing Centre. JURECA - Jülich research on exascale cluster architectures. http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JURECA/JURECA_node.html.
- [24] Jülich Supercomputing Centre. JUWELS - Jülich wizard for european leadership science. http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JUWELS/JUWELS_news.html.
- [25] M. Lüscher. Solution of the Dirac Equation in Lattice QCD Using a Domain Decomposition Method. *Comput. Phys. Commun.*, 156:209–220, 2004.
- [26] M. Lüscher. Local coherence and deflation of the low quark modes in lattice QCD. *JHEP*, 07(2007)081, 2007.
- [27] I. Montvay and G. Münster. *Quantum Fields on a Lattice*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 1994.
- [28] R. B. Morgan and D. S. Scott. Generalizations of Davidson’s method for computing eigenvalues of sparse symmetric matrices. *SIAM J. Sci. Stat. Comput.*, 7(3):817–825, 1986.
- [29] H. Neff, N. Eicker, T. Lippert, J. W. Negele, and K. Schilling. On the low Fermionic Eigenmode Dominance in QCD on the Lattice. *Phys. Rev.*, D64:114509, 2001.
- [30] Y. Notay. Combination of Jacobi-Davidson and conjugate gradients for the partial symmetric eigenproblem. *Numer. Linear Algebra Appl.*, 9(1):21–44, 2002.
- [31] J. C. Osborn, R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen, and C. Rebbi. Multigrid solver for clover fermions. *PoS, LATTICE2010:037*, 2010.
- [32] C. C. Paige, B. N. Parlett, and H. A. van der Vorst. Approximate solutions and eigenvalue bounds from Krylov subspaces. *Numer. Linear Algebra Appl.*, 2(2):115–133, 1998.
- [33] Y. Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM J. Sci. Comput.*, 14(2):461–469, 1992.

- [34] Y. Saad. *Numerical Methods for Large Eigenvalue Problems*. SIAM, Philadelphia, PA, USA, 2nd edition, 2011.
- [35] B. Sheikholeslami and R. Wohlert. Improved Continuum Limit Lattice Action for QCD With Wilson Fermions. *Nucl. Phys.*, B259:572, 1985.
- [36] G. L. G. Sleijpen and H. A. van der Vorst. A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 17(2):401–425, 1996.
- [37] G. D. Smith. *Numerical Solution of Partial Differential Equations: Finite Difference Methods*. Oxford Applied Mathematics and Computing Science Series. Clarendon Press, 1985.
- [38] D.C. Sorensen, R.B. Lehoucq, C. Yang, and K. Maschhoff. PARPACK. <http://www.caam.rice.edu/software/ARPACK>, used version: 2.1, September 1996.
- [39] A. Stathopoulos. Nearly optimal preconditioned methods for Hermitian eigenproblems under limited memory. part i: Seeking one eigenvalue. *SIAM J. Sci. Comput.*, 29(2):481–514, 2007.
- [40] A. Stathopoulos and J. R. McCombs. Nearly optimal preconditioned methods for Hermitian eigenproblems under limited memory. part ii: Seeking many eigenvalues. *SIAM J. Sci. Comput.*, 29(5):2162–2188, 2007.
- [41] A. Stathopoulos and J. R. McCombs. PRIMME: PReconditioned Iterative MultiMethod Eigensolver: Methods and software description. *ACM Transactions on Mathematical Software*, 37(2):21:1–21:30, 2010.
- [42] C. Vuik and H. van der Vorst. GMRESR: A family of nested GMRES methods. *Numerical Linear Algebra with Applications*, 1(4):369–386, 1994.
- [43] L. Wu, E. Romero, and A. Stathopoulos. PRIMME_SVDS: A high-performance preconditioned SVD solver for accurate large-scale computations. *SIAM J. Sci. Comput.*, to appear, abs/1607.01404, 2016.