

Bergische Universität Wuppertal

Fachbereich Mathematik und Naturwissenschaften

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

Preprint BUW-IMACM 11/31

Andreas Frommer, Karsten Kahl, Stefan Krieg, Björn Leder, Matthias Rottmann

Aggregation-based Multilevel Methods for Lattice QCD

December 2011

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

Aggregation-based Multilevel Methods for Lattice QCD

A. Frommer

Department of Mathematics, Bergische Universität Wuppertal, 42097 Germany

E-mail: frommer@math.uni-wuppertal.de

K. Kahl

Department of Mathematics, Bergische Universität Wuppertal, 42097 Germany

E-mail: kkahl@math.uni-wuppertal.de

S. Krieg

Institute for Advanced Simulation, Jülich Supercomputing Centre,

Forschungszentrum Jülich, 52428 Germany

E-mail: s.krieg@fz-juelich.de

B. Leder

Department of Mathematics, Bergische Universität Wuppertal, 42097 Germany

E-mail: leder@math.uni-wuppertal.de

M. Rottmann*

Department of Mathematics, Bergische Universität Wuppertal, 42097 Germany

E-mail: rottmann@math.uni-wuppertal.de

In Lattice QCD computations a substantial amount of work is spent in solving the Dirac equation. In the recent past it has been observed that conventional Krylov solvers tend to critically slow down for large lattices and small quark masses. We present a Schwarz alternating procedure (SAP) multilevel method as a solver for the Clover improved Wilson discretization of the Dirac equation. This approach combines two components (SAP and algebraic multigrid) that have separately been used in lattice QCD before. In combination with a bootstrap setup procedure we show that considerable speed-up over conventional Krylov subspace methods for realistic configurations can be achieved.

The XXIX International Symposium on Lattice Field Theory - Lattice 2011

July 10-16, 2011

Squaw Valley, Lake Tahoe, California

*Speaker.

1. Introduction

The most costly task in lattice QCD computations is the solution of large sparse linear systems of equations

$$Dz = b, \quad (1.1)$$

where D is a discretization of the Dirac operator. Here we consider the Wilson discretization $D = D(U) + m$ which couples only nearest neighbors and depends on a gauge field U and a mass parameter m . Usually z is calculated by a Krylov subspace method (e.g. CGN, GCR, BiCGStab). Those methods suffer from critical slowing down when approaching the critical mass as well as lattice spacing $a = 0$. Thus it is of utmost importance to develop preconditioners for these methods that remedy these scaling problems.

In the recent past preconditioners based on domain decomposition (DD) for the solution of (1.1) have been proposed in [9]. Although DD methods excel in supercomputing environments due to their high inherent parallelism they are unable to remedy the scaling problems completely unless they are combined with a multilevel approach. Thus we combine the DD approach with an algebraic multigrid hierarchy based on a bootstrap aggregation framework [3, 5]. Our approach is similar in construction to the one introduced in [1, 2, 4] where it has been shown that using such algebraic multigrid approaches can remedy the scaling problems in QCD computations. As in [1, 2, 4] we obtain a multilevel hierarchy using non-smoothed aggregation. The difference is that we replace the multigrid smoother by a DD approach, expecting a gain in efficiency on highly parallel machines.

In section 2 we introduce the concept of DD methods before we explain the construction of our algebraic multilevel method in some detail in section 3. Thereafter we give numerical results of our method in section 4 and finish with some concluding remarks.

2. Domain Decomposition

Domain decomposition methods were developed as iterative solvers for linear systems arising from discretizations of PDEs. The main idea consists of solving the system on the whole domain by repeatedly solving smaller systems with less degrees of freedom on smaller subdomains.

Consider a block decomposition $\{\mathcal{L}_i : i = 1, \dots, k\}$ of a lattice \mathcal{L} (Figure 1 illustrates a 2D example). The corresponding trivial embeddings and block solvers are denoted by $I_{\mathcal{L}_i} : \mathcal{L}_i \rightarrow \mathcal{L}$ and $B_i = I_{\mathcal{L}_i} [I_{\mathcal{L}_i}^T D I_{\mathcal{L}_i}]^{-1} I_{\mathcal{L}_i}^T$. Note that the trivial embedding $I_{\mathcal{L}_i}$ is just the restriction of the identity on \mathcal{L} to \mathcal{L}_i . Then one iteration of a domain decomposition method consists of solving each of the block systems $e \leftarrow B_i r$, interleaved with a number of residual updates $r = b - Dz$. In the two extreme cases where one does only one residual update before solving all block systems or when the residual is updated after each block solution, the corresponding error propagators are given by

$$1 - \left(\sum_{i=1}^k B_i \right) A \quad \text{and} \quad \prod_{i=1}^k (1 - B_i A). \quad (2.1)$$

These methods go back to H. Schwarz [11] and thus are called additive Schwarz and multiplicative Schwarz method, respectively. The block systems of the additive variant can be solved simultaneously while the multiplicative variant is inherently sequential. Its advantage is that it spreads the information faster on the lattice as a solution of a block system uses previous solutions of other block systems.

Algorithm 2.1 Red-Black Schwarz

```

1: for  $c = 1$  to 2 do
2:    $r \leftarrow b - Dz$ 
3:   for all  $i \in \{1, \dots, k\}$  with  $color(i) = c$  do
4:      $z \leftarrow z + B_i r$ 
5:   end for
6: end for

```

The two methods can be combined to exploit advantages of both methods. Coloring the blocks such that no adjacent blocks have the same color, a residual update on one block no longer influences the residual on blocks of the same color. Thus it suffices to perform the update once for each color. All blocks of the same color can then be computed simultaneously as described in Algorithm 2.1. Such a DD approach has been applied to solve (1.1) in [9], where it has been named Schwarz Alternating Procedure (SAP).

Typically the solution of the block system $e = B_i r$ is approximated by a few iterations of a Krylov subspace method (e.g. GMRES), and the DD method itself is in turn used as a preconditioner for a (flexible) Krylov subspace method. As illustrated in Figure 2 we observe that SAP is able to reduce error components belonging to a large part of the spectrum very well but a small part belonging to eigenvectors (EVs) to small eigenvalues (EWs) remains intractable. For larger configurations the number of EWs with small magnitude of the Dirac operator gets larger, which yields an explanation why SAP is not able to remedy the scaling problem as the number of intractable eigenvectors increases as well. Though, the seen behavior of damping large EVs is desirable for an iterative method to be used as a smoother in a multigrid method and motivated us to use it in this context.

3. Algebraic Multigrid

A multigrid method typically consists of a simple iterative method called smoother and complementary coarse grid correction. As motivated in section 2 we deem SAP suitable for the use as a smoother since it is cheap to compute and reduces the error efficiently on a large part of the spectrum. The main idea of multigrid is to treat the error that is left after a few iterations of the




Figure 1: block decomposed lattice (reduced to 2D) with 2 colors




Figure 2: error component reduction in terms of EVs on a 4^4 lattice with 2^4 blocks, EWs sorted by magnitude

smoother within a smaller subspace where the troublesome error components can be approximated. More precisely we want to define and solve a “coarse” linear system

$$D_c z_c = b_c \quad (3.1)$$

with a much smaller operator $D_c : \mathbb{C}^{n_c} \rightarrow \mathbb{C}^{n_c}$ in order to reduce the error on the critical part of (1.1). To this purpose we have to define linear maps $R : \mathbb{C}^n \rightarrow \mathbb{C}^{n_c}$ to restrict information based on the current residual $r = b - Dz$ to the subspace and a linear map $P : \mathbb{C}^{n_c} \rightarrow \mathbb{C}^n$ to interpolate the information that we obtain from solving (3.1) back to \mathbb{C}^n where (1.1) is given. This yields a subspace correction

$$z \leftarrow z + PD_c^{-1} Rr \quad (3.2)$$

with the corresponding error propagator $1 - PD_c^{-1} RD$. As D_c should resemble the action of D on the troublesome subspace approximated by $\text{span}(P)$, the action of D_c is chosen as the action of D on interpolated vectors which are restricted afterwards. Formally this amounts to a Petrov-Galerkin formulation of the coarse operator as $D_c = RDP$. With this choice of D_c the error propagator of the subspace correction is given by $1 - P(RDP)^{-1} RD$. In order to benefit from such a subspace correction, solving (3.1) has to be much cheaper than solving (1.1). That is n_c should be small compared to n , and D_c should be sparse. As the dimension of the troublesome subspace grows with n (cf. [10]) we do not want to fix n_c (like in deflation methods) but want to find a sparse description of D_c on that subspace.

Once D_c is found a basic two level algorithm consists of the alternating application of smoother and subspace correction. This procedure can be recursively extended by formulating a two level algorithm of this kind for the computation of (3.1) until we get an operator which is small enough to solve (3.1) directly.

Aggregation Based Interpolation: We decided to adjust an aggregation based interpolation, that in turn yields the subspace correction, as the complementary component to the SAP smoother. Due to the fact that the coarse grid correction in (3.2) only acts on error components in $\text{range}(P)$, it should approximate the subspace spanned by eigenvectors to eigenvalues of small magnitude of D (cf. Figure 2). In the multigrid literature, P is built from right, and R from left EVs corresponding to EWs of small magnitude of D . Due to the spectral properties of the Wilson Dirac operator it is natural to choose $R = \gamma_5 P$. Furthermore, with additional assumptions on its structure we can choose $R = P^\dagger$ according to [2]. Therefore we define the aggregates in such a way that there exists γ_5^ξ such that $\gamma_5 P = P \gamma_5^\xi$. Note, that with these assumptions on P the error propagator of the subspace correction (3.2) is given by $z \leftarrow z + P(P^\dagger DP)^{-1} P^\dagger r$.

With this structure of P in mind, we define its entries based on a set of test vectors $\{v_1, \dots, v_N\}$ whose span approximates the troublesome subspace and a set of aggregates $\{\mathcal{A}_1, \dots, \mathcal{A}_s\}$. The aggregates can be realized as another block decomposition of the lattice. Note that the DD smoother and the interpolation do not have to share a common block decomposition. The interpolation P is then given by decomposing the test vectors over the aggregates (cf. Figure 3). Hence P is a linear map from the coarse grid to the fine grid, defined by

$$Pe_j := I_{\mathcal{A}_j}^T v_{((j-1) \bmod N)+1} \quad \text{for } j = 1, \dots, N \cdot s$$

$$(v_1, \dots, v_N) = \left(\begin{array}{c} | \\ | \\ | \\ | \end{array} \right) \mapsto P = \left(\begin{array}{c} \boxed{} \\ \boxed{} \\ \vdots \\ \boxed{} \end{array} \right)$$

Figure 3: construction of the interpolation (operator based point of view)

where e_j is the j -th unit vector. In order to have $P^\dagger P = I$, the test vectors are locally orthonormalized over the aggregates. Note that due to the aggregation structure of P and R the sparsity/connection structure of D_c resembles the one of D , i.e., the corresponding graphs of the operators are regular four dimensional grids. Thus we can apply (3.2) recursively to (3.1) and obtain a hierarchy of coarser grids and coarser operators. This construction of P is similar to constructions found in [1, 2, 4, 10]. More precisely, the structure of the interpolation operators is identical but the test vectors v_i used to build them and thus the actions of the operators are different.

Bootstrap Setup: A critical part of the construction of an efficient multigrid hierarchy is the computation of the test vectors used in the definition of P . The setup procedure we employ for this task is divided into two parts. We start with a set of random test vectors v_1, \dots, v_N and apply a small number of smoother iterations S^V to them. During this procedure the smoothed test vectors $S^V v_i$ are kept orthonormal. After that a temporary operator on the next coarser grid is computed according to the aggregation based construction. This procedure is continued recursively on the coarser grids until we get an operator on the coarsest grid. Herein we restrict smoothed test vectors with P^\dagger to the next coarser grid. In what follows we omit the level indices for the sake of simplicity.

Since the subspace $\text{range}(P)$ should contain a significant part of the eigenvectors to eigenvalues of small magnitude of the fine grid space \mathbb{C}^n , the EWs of small magnitude of $D_c = P^\dagger D P$ should approximate those of D . For the corresponding EVs the relations $P^\dagger P = 1$ and $P^\dagger (D P \varphi - \lambda P \varphi) = 0$ imply $D P \varphi \approx \lambda P \varphi$. Thus the second part of our setup procedure starts with the calculation of N approximations to the smallest EVs and EWs $\{(\lambda_i, \varphi_i)\}$ of $P^\dagger D P$ by means of harmonic Ritz vectors and values. The approximate EVs φ_i for $i = 1, \dots, N$ are successively interpolated to the next finer grid and smoothed towards their harmonic Ritz values with some steps of SAP with iteration matrix $S(\lambda_i)^l$. In this case the local inverses for the pair (λ_i, φ_i) are given by

$$B_j(\lambda_i) = I_{\mathcal{L}_j} [I_{\mathcal{L}_j}^T (D - \lambda_i) I_{\mathcal{L}_j}]^{-1} I_{\mathcal{L}_j}^T.$$

The resulting set of vectors and the old test vectors $V := \{S(\lambda_i)^l P \varphi_i : i = 1, \dots, N\} \cup \{S^V v_j : j = 1, \dots, N\}$ are again reduced to N vectors. In order to preserve the most significant information, the N smallest singular values and their corresponding vectors of the $2N \times 2N$ matrix $(DV)^\dagger DV$ are calculated and the corresponding N linear combinations of the vectors of V are the final vectors \hat{V} . They define the interpolation from the next coarser grid to the current grid and the operator on the next coarser grid. The second part of the setup executes this procedure exactly once, starting on the coarsest grid.

Figures 4 and 5 illustrate the influence of the second part of the setup procedure on the coarse grid operator in a two level hierarchy. The smallest eigenvalues of the fine grid operator are much better represented on the coarse grid. For other setup procedures, see e.g. [1, 3, 5].




Figure 4: Spectra after the first setup step, 4^4 Wilson, 2^4 aggregate size




Figure 5: Spectra after the second setup step, 4^4 Wilson, 2^4 aggregate size

4. Results

Our adaptive DD multilevel solver (α MG-DD) has been implemented in C using the parallelization interface of MP I. Krylov subspace methods have been implemented within a common framework for a fair comparison. For the numerical experiments we have combined our multigrid solver against CGN, i.e. CG on the normal equation. All results were produced with a two level method. The stopping criterion was to reduce the initial residual norm by a factor of at least 10^{10} . The solutions of the block systems within SAP were approximated by 3 iterations of GMRES and the coarse grid equation (3.1) was approximately solved with 12 iterations of GMRES. For the outer flexible GMRES routine we chose a restart length of 25. All results have been computed on Juropa at the Jülich Supercomputing Centre.

		FGMRES + α MG-DD	CGN
setup	timings	45.22s	-
solve	iterations	9	4476
	timings	6.97s	121.82s
total	timings	52.19s	121.82s

Table 1: 48×24^3 Clover Wilson-Dirac, $\beta = 5.3$ ($1/a = 2.56$ GeV), $\kappa = 0.13590$ ($m_\pi = 630$ MeV), $c_{sw} = 1.90952$, generated using public code with parameters from L. Del Debbio [6], block-size 3^4 , coarsening $3^4 \times 6 \rightarrow 20$, 512 cores (Juropa at JSC).

		FGMRES + α MG-DD	CGN
setup	timings	24.85s	-
solve	iterations	17	9181
	timings	5.05s	119.04s
total	timings	29.90s	119.04s

Table 2: 64×32^3 Wilson-Dirac 2HEX smeared tree level improved Clover, $\beta = 3.5$ ($1/a = 2.130$ GeV), $\kappa = 0.12646$ ($m_\pi = 300$ MeV), $c_{sw} = 1.0$, provided by the BMW collaboration [7, 8], block-size 2^4 , coarsening $4^4 \times 6 \rightarrow 20$, 4096 cores (Juropa at JSC).

In the cases shown our multigrid solver was 17 and 24 times faster than CGN. Including the setup time we are still twice and four times as fast as CGN. Since the setup has to be done only once the benefits of our approach are larger the more right hand sides there are to be solved.

5. Summary and Outlook

The developed method combining DD techniques and algebraic multigrid shows great potential to speed-up calculations of propagators in lattice QCD. Even for single right hand sides our method outperforms conventional Krylov subspace methods with the potential of an even more significant speed-up when solving for many right hand sides. This result is mainly due to the introduction of the highly parallel DD smoother and the bootstrap setup into the algebraic multigrid

method. While the first speeds up the setup of the method *and* the subsequent solution the latter significantly speeds up the setup of the multigrid hierarchy, which in general is a bottleneck in algebraic multigrid methods especially compared to setup-free Krylov subspace methods. We are currently working on an optimized version of the code that should be able to run on large-scale parallel machines and extend our testing of the method towards larger lattices and lighter quark masses. In the near future we plan to incorporate our algorithm into the production codes of our collaborators within SFB TR55.

Acknowledgments: This work is funded by Deutsche Forschungsgemeinschaft (DFG) Transregional Collaborative Research Centre 55 (SFB TR55).

References

- [1] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen, J. C. Osborn, and C. Rebbi, *Multigrid solver for clover fermions*, The XXVIII International Symposium on Lattice Field Theory, 2010, arXiv:1011.2775v1 [hep-lat].
- [2] 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), arXiv:1005.3043v2 [hep-lat].
- [3] A. Brandt, J. Brannick, K. Kahl, and I. Livshits, *Bootstrap AMG.*, SIAM J. Sci. Comput. **33** (2011), no. 2, 612–632.
- [4] 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), arXiv:0707.4018v1 [hep-lat].
- [5] M. Brezina, R. Falgout, S. MacLachlan, T. Manteuffel, S. McCormick, and J. Ruge, *Adaptive smoothed aggregation (α SA) multigrid*, SIAM Review 47, No. 2 (2005), 317–346.
- [6] L. Del Debbio, Leonardo Giusti, M. Lüscher, R. Petronzio, and N. Tantalo, *QCD with light Wilson quarks on fine lattices (I): First experiences and physics results*, JHEP **0702** (2007), 056.
- [7] S. Durr, Z. Fodor, C. Hoelbling, S. D. Katz, S. Krieg, T. Kurth, L. Lellouch, T. Lippert, K. K. Szabo, and G. Vulvert, *Lattice QCD at the physical point: Simulation and analysis details*, arXiv:1011.2711v1 [hep-lat].
- [8] ———, *Lattice QCD at the physical point: light quark masses*, Phys. Lett. B701 (2011), 265–268, arXiv:1011.2403v2 [hep-lat].
- [9] Martin Lüscher, *Solution of the Dirac equation in lattice QCD using a domain decomposition method*, Comput. Phys. Commun. 156 (2004), 209–220, arXiv:hep-lat/0310048v1.
- [10] ———, *Local coherence and deflation of the low quark modes in lattice QCD*, JHEP0707:081 (2007), arXiv:0706.2298v4 [hep-lat].
- [11] Hermann Schwarz, *Gesammelte mathematische Abhandlungen*, Vierteljahrsschrift Naturforsch. Ges. Zürich (1870), 272–286.