

Bergische Universität Wuppertal

Fachbereich Mathematik und Naturwissenschaften

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

Preprint BUW-IMACM 12/11

Andreas Frommer, Andrea Nobile and Paul Zingler

**Deflation and Flexible SAP-Preconditioning of
GMRES in Lattice QCD Simulation**

April 2012

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

Deflation and Flexible SAP-Preconditioning of GMRES in Lattice QCD Simulation

Andreas Frommer^b, Andrea Nobile^a, Paul Zingler^b

^a*Institute for Theoretical Physics, University of Regensburg, 93040 Regensburg, Germany*

^b*Department of Mathematics, University of Wuppertal, 42097 Wuppertal, Germany*

Abstract

The simulation of lattice QCD on massively parallel computers stimulated the development of scalable algorithms for the solution of sparse linear systems. We tackle the problem of the Wilson-Dirac operator inversion by combining a Schwarz alternating procedure (SAP) in multiplicative form with a flexible variant of the GMRES-DR algorithm. We show that restarted GMRES is not able to converge when the system is poorly conditioned. By adding deflation in the form of the FGMRES-DR algorithm, an important fraction of the information produced by the iterates is kept between successive restarts leading to convergence in cases in which FGMRES stagnates.

1. Introduction

The simulation of Lattice QCD is among one of the most challenging problems in computational science because of its enormous computational cost, see [1], e.g.. With sufficiently powerful computers becoming available during the last fifteen years, the simulation of Lattice QCD including dynamical fermions became a reality. The de-facto standard algorithm used to accomplish this task is the Hybrid Monte Carlo (HMC) algorithm [2]. HMC introduces a fictitious time, in which a dynamical system is evolved according to the system Hamiltonian. In order to integrate the Hamiltonian equations of motion, the forces due to both, gauge field and pseudo-fermion field, have to be evaluated accurately. The by far most demanding numerical task is the inversion of the fermionic Dirac matrix, needed to evaluate the fermion force. In this paper we investigate a class of inverters for the fermionic Dirac matrix which is particularly well adapted to current high performance hardware architecture. Our work was primarily inspired by the architecture of the QPACE machine, see [3, 4], but it is of a general nature and applies to all architectures where data movements rather than arithmetic operations tend to be the limiting factor on performance.

^{*}Supported by DFG collaborative research center SFB/TR-55 “Hadron Physics from Lattice QCD”.

The development of machine-aware algorithms has become particularly important during the last years as computer architectures move towards dramatically increased floating point (FP) speed and increased on-chip parallelism. For reasons of cost and because of physical limitations, these improvements are not accompanied by comparable progresses in the bandwidths and latencies of main memory and processor interconnects. It is predicted that this situation will be even aggravated as we move towards the Exascale area, see [5], e.g..

An often used rough indicator of balance between processor and memory subsystems is the ratio between peak main memory bandwidth and peak floating point performance. In the past few years this value went down from ≈ 1 to $\approx 1/8$ and less for the CELL processor, modern GPUs and general purpose x86 systems. This increasing gap is the main reason for a line of computer science research on software techniques aimed at maximizing cache utilization. A famous example is [6].

The algorithm cost evaluation must take into account not only the FP operations (as it was traditionally done some years ago when FP was the most relevant bottleneck in numerical applications), but has to carefully evaluate the amount and pattern of data movement among the levels of memory hierarchy and the network as these are now the most limiting factors.

Algorithms that are more suitable for these increasingly unbalanced computer architectures are particularly attractive as they can give easier access to the FP performances offered today.

We describe an algorithm for the Dirac-Wilson matrix inversion that is able to exploit the architectural features of modern machines. We show the convergence behaviour and demonstrate the performance with our implementation on the peculiar SFB-TR55 QPACE architecture based on the IBM PowerXCell8i processor with state of the art lattices.

2. Domain decomposition and SAP

The Wilson fermion matrix D_w [7] describes a (periodic) nearest neighbor coupling on the 4-dimensional space-time lattice. Without further explanation, our desire is to solve the linear system

$$D_w x = b.$$

Owed to the complexity and size of D_w , a direct solution is not feasible and we therefore exploit the structure of D_w . We decompose the overall lattice into sublattices, called *domains*. For each domain, we obtain a local Wilson fermion operator by simple restriction; at the local boundary of a domain, we remove the couplings with neighboring domains. To describe this formally, let $n = (n_1, n_2, n_3, n_4)$ stand for a 4-tupel describing a lattice site and let $L = \{n : (0, 0, 0, 0) \leq n \leq (d_1 - 1, d_2 - 1, d_3 - 1, d_4 - 1)\}$ denote the overall lattice. The

Dirac Wilson operator can then be written as

$$[D_w]_{nm} = \delta_{n,m-\kappa} \left(\sum_{j=1}^4 (1 + \gamma_j) U_{n,j} \delta_{n+\hat{j},m} + (1 - \gamma_j) U_{n-\hat{j},j}^\dagger \delta_{n-\hat{j},m} \right), \quad n, m \in L, \quad (1)$$

where $n \pm \hat{j}$ is to be taken mod d_j .

A domain L^ℓ is characterized by the bounds a_j^ℓ, d_j^ℓ for its lattice sites which are given by $L^\ell = \{n : (a_1^\ell, a_2^\ell, a_3^\ell, a_4^\ell) \leq n \leq (d_1^\ell - 1, d_2^\ell - 1, d_3^\ell - 1, d_4^\ell - 1)\}$. The local Dirac operator for domain L^ℓ is thus given as

$$[D_w^\ell]_{nm} = \delta_{n,m-\kappa} \left(\sum_{j=1}^4 {}' (1 + \gamma_j) U_{n,j} \delta_{n+\hat{j},m} + (1 - \gamma_j) U_{n-\hat{j},j}^\dagger \delta_{n-\hat{j},m} \right), \quad n, m \in L^\ell, \quad (2)$$

where \sum' indicates that the sum is to be taken only over those contributions for which $n + \hat{j}$ or $n - \hat{j}$ is in L^ℓ .

Assuming that the local Wilson operators are easy to invert, we can perform the following basic domain decomposition iteration, known as the multiplicative Schwarz method in the domain decomposition literature [8, 9].

Algorithm 1: Multiplicative Schwarz method

```

choose initial guess  $x$ 
repeat
  for  $\ell = 1, \dots, p$  do
    compute residual  $r^\ell$  for domain  $\ell$ :  $r^\ell = (b - D_w x)^\ell$ 
    solve  $D_w^\ell y^\ell = r^\ell$ 
    update: replace part  $x^\ell$  in  $x$  by  $y^\ell$ 
  until until convergence

```

Herein, the superscript ℓ indicates that we only take the part which corresponds to the subdomain ℓ . Due to the nearest neighbor coupling of D_w , computing r^ℓ is a local operation which involves only those components of x which belong to subdomain L^ℓ , i.e. x^ℓ and those from neighboring subdomains which are at the boundary to domain L^ℓ .

If we use a red-black coloring of the subdomains as depicted in Fig. 1 for a two-dimensional lattice, and we have an even number of sub-domains in each dimension of the lattice, then there is no coupling between subdomains of the same color. When we order all red (say) domains first, we see that in the algorithm above the values of the residual r^ℓ does not depend on whether or not we updated the part x^j of the iterate for the other subdomains of the same color. If we proceed by colors, we end up with the following variant of the basic domain decomposition iteration which is termed the Schwarz alternating procedure (SAP) in [10].




Figure 1: domain decomposition in 2 dimensions

Algorithm 2: SAP iteration

```

choose initial guess  $x$ 
repeat
  for all red domains  $\ell$  do
    compute residual  $r^\ell$  for domain  $\ell$ :  $r^\ell = (b - D_w x)^\ell$ 
    solve  $D_w^\ell y^\ell = r^\ell$ 
  update: replace part  $x^\ell$  in  $x$  by  $y^\ell$  for all red domains  $\ell$ 
  for all black domains  $\ell$  do
    compute residual  $r^\ell$  for domain  $\ell$ :  $r^\ell = (b - D_w x)^\ell$ 
    solve  $D_w^\ell y^\ell = r^\ell$ 
  update: replace part  $x^\ell$  in  $x$  by  $y^\ell$  for all black domains  $\ell$ 
until until convergence

```

3. SAP in computational practice

In computational practice, the SAP method can rarely be used directly in the form given by Algorithm 2. On one hand it is usually too costly to solve the local systems $D_w^\ell y^\ell = r^\ell$ exactly. One rather computes an approximation to y^ℓ using another, *inner* iterative method on the systems $D_w^\ell \tilde{y}^\ell = r^\ell$. Usually, the fastest overall method arises if one requires a fairly low accuracy in the inner iteration. For example, one requires that the approximate solution \tilde{y}^ℓ reduces the initial residual by just a factor of 10, i.e. one requires

$$\|r^\ell - D_w^\ell \tilde{y}^\ell\| \leq 0.1 \cdot \|r^\ell\|.$$

Such a method is termed an *inexact* SAP method.

On the other hand, even the exact SAP iteration may diverge for the Wilson Dirac operator when the hopping parameter κ approaches the critical value κ_c . Indeed, most of the convergence theory for multiplicative Schwarz methods assumes that the operator is hermitian and positive definite, see [8, 9], which is not the case for the Wilson Dirac operator.

The usage of SAP methods is particularly suitable on parallel machines and in general on machines with a memory hierarchy. Provided that at least two subdomains are assigned to a machine node, there is no need for network communication during the solution of the Wilson equation on the subdomains, since all the necessary data is on the node. Data only needs to be communicated for the update of the residual. In practice, there are more than two subdomains per node giving the possibility to overlap network communication and computation. The required network bandwidth, and in particular latency, is less stringent than for a simple Krylov solver, since the approximate solution of the subdomain Wilson equation requires some iterations. The required machine performance on global network operations (i.e. global sum, barrier, etc.) is also typically reduced, since the frequency of such operations is lower than in the case of a Krylov solver. We have chosen SAP as a preconditioner to be used on the most time-critical solves on our machine QPACE. The QPACE torus network [4] is designed with a typical Krylov solver in mind and has therefore high bandwidth and low latency, nevertheless SAP is particularly appealing due to the peculiar memory hierarchy architecture of the IBM PowerXCell 8i Processor. On this processor, the fast on-chip memory is not managed automatically by the hardware as on standard processors (cache), but its control is completely left to the programmer. If the size of the subdomain is chosen such that it fits the fast memory, no main memory access is necessary for the subdomain solve, giving the possibility to achieve impressive sustained performance for the subdomain solves (up to 50% of the peak) [11], [12]. This feature of the SAP algorithm is interesting also for future architectures as the main memory performance will continue to improve at a lower rate than processor performance.

4. Domain decomposition preconditioner

Exact or inexact SAP iterations can be used as a preconditioner in a Krylov subspace method like GMRES [13] or GCR [14]. This will usually result in an iterative process which converges faster than the Krylov subspace method without preconditioning or the stand-alone SAP iteration. We will discuss GMRES and GCR in quite some detail later in section 5. At this point, we concentrate on aspects pertaining to the use of *inexact* SAP iterations in the context of preconditioning.

Loosely speaking, the speed of convergence of a Krylov subspace method for a generic system

$$Ax = b$$

gets faster the closer A is to the identity. So, to accelerate the convergence one tries to construct (or implicitly use) a matrix M , the preconditioner, such that the matrix AM of the (right) preconditioned system

$$(AM)y = b, \quad x = My$$

is close to the identity. In this sense, M should be an approximation to the inverse of A . Such M is given implicitly if we perform some steps of an iterative procedure like SAP.

Each step of the Krylov subspace method requires one multiplication with A and one with M . If we perform exact SAP with a fixed number, s say, of iterative steps, the preconditioner M can be expressed as the truncated series

$$M = \sum_{j=0}^{s-1} (A_L^{-1} A_U) A_L^{-1},$$

where the matrices A_L and A_U result from a splitting $A = A_L - A_U$ of the matrix A . To be specific, for the SAP iteration from Algorithm 2, where $A = D_w$, this splitting is based on the red-black permuted matrix D_w given as

$$D_w = \begin{bmatrix} D_w^{bb} & D_w^{br} \\ D_w^{rb} & D_w^{rr} \end{bmatrix}.$$

Herein, the blocks D_w^{bb} and D_w^{rr} are made up of all the local Dirac operators D_w^ℓ of the black and the red subdomains, respectively and the off-diagonal blocks contain the couplings between subdomains of different color. Then

$$A_L = \begin{bmatrix} D_w^{bb} & 0 \\ D_w^{rb} & D_w^{rr} \end{bmatrix}, \quad A_U = \begin{bmatrix} 0 & -D_w^{br} \\ 0 & 0 \end{bmatrix}.$$

In *inexact* SAP we invert A_L only approximately, using some inner iteration on the diagonal blocks D_w^ℓ . If this inner iteration is non-stationary, as it is the case if we perform some steps of GMRES or MR [15], the implicitly given preconditioner changes at each iteration. This is a fact one has to account for when formulating the preconditioned Krylov subspace methods. In the linear

algebra literature such methods are termed *flexible*, see [16, 15] since they adapt themselves to situations where the preconditioner varies from one iteration to the next.

The next chapter will first introduce the flexible restarted GMRES method before turning to the crucial subject of this paper, namely its acceleration via deflation.

5. F-GMRES-DR

To simplify the presentation we will use a generic notation in this chapter, i.e. the linear system to solve will be denoted

$$Ax = b.$$

A preconditioner will be denoted M . In a variable preconditioning context, where the preconditioner will depend on the current iterative step j , we denote the preconditioner by M_j . Note that M_j will usually not be given explicitly but rather be the result of some steps of a non-stationary iteration like inexact SAP.

Particular aspects for (inexact) SAP preconditioning of the Wilson Dirac operator will be discussed later.

5.1. GMRES and GCR

We consider a generic non-singular linear system

$$Ax = b \text{ with solution } x_* = A^{-1}b.$$

Given an initial guess x_0 and its residual $r_0 = b - Ax_0$, the j -th Krylov subspace $K_j(A, r_0)$ is defined as

$$K_j(A, r_0) = \text{span}\{r_0, Ar_0, \dots, A^{j-1}r_0\}.$$

For ease of notation, we will often write K_j instead of $K_j(A, r_0)$. It is known that the solution $x_* = A^{-1}b$ is contained in the space $x_0 + K_{j_*}$ where j_* denotes the smallest j for which $K_j = K_{j+1}$. Krylov subspace methods work by iteratively choosing approximations x_j from $x_0 + K_j$ to approximate x_* . The GMRES and GCR method are Krylov subspace methods which both obtain the “optimal” x_j in the sense that the residual $r_j = b - Ax_j$ satisfies the minimal residual condition

$$\|r_j\| = \min_{x \in x_0 + K_j} \|b - Ax\|.$$

Note that for any subspace \mathcal{V} of \mathbb{C}^n we have

$$x = x_0 + \delta x \text{ minimizes } \|b - Ax\| \text{ over } x_0 + \mathcal{V} \Leftrightarrow r_0 - Ax \perp A\mathcal{V}, \quad (3)$$

where $r_0 - Ax = b - A(x_0 + \delta x)$ is the residual of $x = x_0 + \delta x$. The equivalence (3) holds because the minimizer δx is such that $A\delta x$ represents the orthogonal

projection of $b - A(x_0 + \delta x)$ onto AV . So the minimal residual condition for GMRES or GCR is equivalent to

$$r_j \perp AK_j. \quad (4)$$

GMRES works by building a sequence of orthonormal vectors v_1, \dots, v_j which span $K_j(A, r_0)$. The coefficients η_i in $x_j = x_0 + \sum_{i=1}^j \eta_i v_i$ are then obtained by solving a small $(j+1) \times j$ least squares problem. GCR computes a sequence of orthonormal vectors w_1, \dots, w_j which span AK_j , and an auxiliary sequence s_1, \dots, s_j with $w_j = As_j$. The iterate x_j is obtained by an update $x_j = x_{j-1} + \alpha_j s_{j-1}$. Both, GMRES and GCR, suffer from the fact that the number of vectors stored and the cost for the orthogonalizations increase linearly with n . A remedy is to restart or truncate the methods. GCR requires twice as many vectors to store as GMRES, and its total arithmetic costs are slightly higher than in GMRES. If convergence is slow, GMRES is significantly more stable numerically than GCR as was proven in [17]. There is one advantage of GCR over GMRES, however: While the GMRES algorithm has to be modified (slightly) to allow for variable preconditioning, GCR adapts itself “automatically” to such a situation. For this reason GCR was used in various contexts of variable preconditioners like, for example, [10]. In the present paper we consider situations where there is an additional need for (inexact) deflation in order to get sufficiently fast convergence. While this can be done elegantly and efficiently within the GMRES framework, there seems to be no such way for GCR. This explains why we focus on GMRES from now on.

5.2. Flexible GMRES

Standard (non-variable) right preconditioning of the equation $Ax = b$ means that instead of looking for an iterate $x_j \in x_0 + K_j(A, r_0)$ we now look for one in the space $x_0 + MK_j(AM, r_0)$, where M stands for the preconditioning matrix. The idea is that with a good preconditioner M (a good approximation to A^{-1}), the “search space” $x_0 + MK_j(AM, r_0)$ contains substantially better approximations to the solution than $x_0 + K_j(A, r_0)$. In right preconditioned GMRES [15] we obtain the j -th iterate x_j as

$$x_j = x_0 + M\delta y_j, \text{ where } \delta y_j = \operatorname{argmin}_{\delta y \in K_j(AM, r_0)} \|b - A(x_0 + M\delta y)\|. \quad (5)$$

Right preconditioned GMRES uses the Arnoldi process to compute an orthonormal basis of $K_j(AM, r_0)$, orthogonalizing AMv^j against all previous vectors v_1, \dots, v_j . The solution of (5) can then be obtained as the solution of a small $(m+1) \times m$ least squares problem. We do not discuss this further here, since right preconditioned GMRES appears as a special case of flexible GMRES [16] which we develop in all its details now.

In a flexible context, the preconditioning matrix depends on the iterative step, i.e. we have a new preconditioning matrix M_j in each step j . In a “flexible” Arnoldi process we thus orthogonalize $AM_j v_j$ against all previous vectors v_1, \dots, v_j . It will turn out useful to also keep track of the preconditioned vectors

Algorithm 3: Flexible Arnoldi Process

Input : $A \in \mathbb{C}^{n \times n}$, $b \in \mathbb{C}^n$, an integer m
Output: $V_{m+1} = [v_1 | \dots | v_{m+1}] \in \mathbb{C}^{n \times (m+1)}$, $Z_m = [z_1 | \dots | z_m] \in \mathbb{C}^{n \times m}$,
 $\hat{H}_m = (h_{i,j}) \in \mathbb{C}^{(m+1) \times m}$

```

1  $\beta = \|b\|$ 
2  $v_1 := \frac{1}{\beta}b$ 
3 for  $j = 1, \dots, m$  do
4    $z_j := M_j v_j$  /* preconditioning of new basis vector */
5    $w := Az_j$ 
6   for  $i = 1, \dots, j$  do /* orthogon. against previous vectors */
7      $h_{ij} := v_i^H w$ 
8      $w = w - h_{ij} v_i$ 
9    $h_{j+1,j} := \|w\|$ 
10   $v_{j+1} := w/h_{j+1,j}$  /* normalization */
```

$z_j = M_j v_j$ for future use when computing (flexible) GMRES iterates. So we end up with the flexible Arnoldi process described as Algorithm 3.

The flexible Arnoldi relation

$$AZ_m = V_{m+1} \hat{H}_m \quad (6)$$

summarizes the computations of Algorithm 3, its j -th column representing

$$h_{j+1,j} v_{j+1} = Az_j - \sum_{i=1}^j h_{ij} v_i.$$

Note that the $(m+1) \times m$ matrix \hat{H}_m is upper Hessenberg, i.e. its elements below the first subdiagonal, h_{ij} for $i > j+1$, are all 0. The columns v_j of V_{m+1} are orthonormal, so we have $V_{m+1}^H V_{m+1} = I \in \mathbb{C}^{(m+1) \times (m+1)}$.

In the non-flexible case, i.e. $M_j = M$ for all j , the orthonormal vectors v_1, \dots, v_j span the Krylov subspace $K_j(AM, r_0)$, and the vectors $z_i = M v_i$, $i = 1, \dots, j$ span the search space $MK_j(AM, r_0)$. In this case we can formulate (6) without Z_m to retrieve the standard Arnoldi relation

$$(AM)V_m = V_{m+1} \hat{H}_m. \quad (7)$$

In the case of a variable preconditioner, there is no immediate notion of a (preconditioned) Krylov subspace, but we can still use the space spanned by the vectors z_i , i.e. $\text{range}(Z_m)$ as a search space. This is the approach taken in flexible GMRES (F-GMRES) [16], where we obtain the iterates x_m by imposing the minimal residual condition for $x_m \in x_0 + \text{range}(Z_m)$. So we have to compute

$$\eta = \operatorname{argmin}_{\xi \in \mathbb{C}^m} \|b - A(x_0 + Z_m \xi)\|. \quad (8)$$

Using the flexible Arnoldi relation (6) and the fact that V_{m+1} has orthonormal columns, we have

$$\|b - A(x_0 + Z_m \xi)\| = \|r_0 - V_{m+1} \widehat{H}_m \xi\| = \|V_{m+1} \beta e_1 - V_{m+1} \widehat{H}_m \xi\| = \|\beta e_1 - \widehat{H}_m \xi\|.$$

So solving the $(m+1) \times m$ least squares problem $\eta = \operatorname{argmin}_{\xi \in \mathbb{C}^m} \|\beta e_1 - \widehat{H}_m \xi\|$ gives the coefficient vector η from which we retrieve the FGMRES iterate x_m as

$$x_m = x_0 + Z_m \eta.$$

We note in passing that the $(m+1) \times m$ least squares problem is usually solved via a QR-factorization of \widehat{H}_m . Due to the upper Hessenberg structure of \widehat{H}_m , this can easily be done in an iterative manner by updating the QR-decomposition of \widehat{H}_j to that of \widehat{H}_{j+1} using one additional Givens rotation. Moreover, it is possible to very cheaply update the norm of the residual of the j -th GMRES iterate without explicitly computing the iterate. This can be used to implement an early termination criterion in cases where the j -th (F)GMRES iterate is already accurate enough for $j < m$. Details can be found in [15], e.g.. See also section 5.6.

5.3. Restarts and Deflation

If larger values of m are required to obtain a sufficiently accurate approximation x_m , full F-GMRES as described in the previous section is not feasible since we have to store the $2m$ vectors v_j and z_j and the total cost for the orthogonalizations in the flexible Arnoldi process grow like $\mathcal{O}(nm^2)$. It is therefore common practice to use *restarted* F-GMRES. The integer m is fixed to a moderate value. One then goes through several cycles. Each cycle performs m steps of the full F-GMRES algorithm, with its initial guess given by the approximation to the solution computed in the previous cycle.

While restarts overcome problems with storage and computational complexity, they may severely degrade the convergence behaviour, and it might even happen that the method stagnates at a large residual norm without achieving any further progress. For difficult problems, this phenomenon has been widely observed in the literature, and several cures have been proposed. One of the most appropriate ones is the deflated restart modification which we discuss now.

Upon a restart, all the information contained in the search space \mathcal{K} built up so far is lost. In particular, in the new cycle we do not keep the residuals orthogonal to $A\mathcal{K}$. The idea of deflated restarts is to extract a low-dimensional subspace \mathcal{U} of \mathcal{K} which contains the most important information in order to speed-up convergence, and to at least maintain the residuals orthogonal to $A\mathcal{U}$. The new cycle can then be characterized as an augmenting Krylov subspace method with \mathcal{U} as the augmenting space. The crucial point is that, as we will see in the next section, it is possible to construct this deflating augmenting subspace in such a way that we do not need additional multiplications with A .

5.4. Augmentation by Deflation

We first describe the ideas and technicalities of deflated restarts for the non-flexible case. The extension to FGMRES will be given in section 5.5. So for now we assume we have a fixed preconditioner M in every iteration, and to further simplify the notation we base our whole discussion on the non-preconditioned system $Ax = b$. The (right) preconditioned case is obtained by merely replacing A with AM everywhere.

The presence of eigenvector components corresponding to small eigenvalues in a residual r induces a large contribution of these eigenvectors in the error $e = A^{-1}r$. It is thus desirable to chose the augmenting subspace \mathcal{U} as one that contains approximations to eigenvectors belonging to small eigenvalues. Those can be approximated from the Krylov subspace built up in the previous cycle. If this is done in the right manner, one can establish an Arnoldi type process for the augmented spaces with little additional cost as compared to standard Krylov subspaces, so that one ends up with an efficient method. We now give the details, following [18], which in turn builds on [19, 20, 21]. The key ingredient is to use harmonic Ritz vectors [22].

Definition 1. *Given a matrix $A \in \mathbb{C}^{n \times n}$ and a subspace $\mathcal{K} \subset \mathbb{C}^n$, a harmonic Ritz pair (λ, v) with $\lambda \in \mathbb{C}, v \in \mathcal{K}$ for A with respect to \mathcal{K} satisfies*

$$Av - \lambda v \perp A\mathcal{K}. \quad (9)$$

Harmonic Ritz pairs are well suited to approximate eigenpairs with a small eigenvalue [22]. The following lemma shows how harmonic Ritz pairs can be computed in cases where \mathcal{K} is a Krylov subspace or an appropriately augmented Krylov subspace.

Lemma 1. *Let \mathcal{K} be a subspace of \mathbb{C}^n of dimension m and assume that*

$$A\mathcal{K} \subset \mathcal{K} \oplus \langle s \rangle \text{ for some vector } s \in \mathbb{C}^n.$$

Let the columns v_1, \dots, v_{m+1} of $V_{m+1} \in \mathbb{C}^{n \times (m+1)}$ be orthonormal vectors such that v_1, \dots, v_m span \mathcal{K} and v_1, \dots, v_{m+1} span $\mathcal{K} \oplus \langle s \rangle$, so that there exists a full rank matrix $\hat{H}_m \in \mathbb{C}^{(m+1) \times m}$ with

$$AV_m = V_{m+1}\hat{H}_m. \quad (10)$$

Moreover, let

$$\hat{H}_m = \begin{bmatrix} H_m \\ h_m^H \end{bmatrix} \text{ with } H_m \in \mathbb{C}^{m \times m}, h_m \in \mathbb{C}^m.$$

Then:

- a) *The harmonic Ritz pairs of A w.r.t. \mathcal{K} are given as $(\lambda, V_m g)$, where (λ, g) is an eigenpair of the matrix*

$$H_m + f_m h_m^H, \text{ where } H_m^H f_m = h_m.$$

b) The residual $Au - \lambda u$ of a harmonic Ritz pair satisfies

$$Au - \lambda u \in \langle r \rangle,$$

where r spans the orthogonal complement of AK in $\mathcal{K} \oplus \langle s \rangle$, i.e.

$$AK \oplus_{\perp} \langle r \rangle = \mathcal{K} \oplus \langle s \rangle.$$

Proof. Using (10), the defining property (9) for the Ritz pair (λ, v) with $v = V_m g$ turns into

$$\begin{aligned} & (AV_m)^H(Av - \lambda v) = 0 \\ \Leftrightarrow & (V_{m+1}\widehat{H}_m)^H(V_{m+1}\widehat{H}_m g - \lambda V_m g) = 0 \\ \Leftrightarrow & \widehat{H}_m^H \left(\widehat{H}_m g - \begin{bmatrix} \lambda I_m \\ 0 \end{bmatrix} g \right) = 0. \end{aligned}$$

Since \widehat{H}_m has full rank, the matrix H_m is non-singular. The last equality can thus be further reduced to

$$\begin{aligned} & (H_m^H H_m + h_m h_m^H) g - \lambda H_m^H g = 0 \\ \Leftrightarrow & (H_m + f_m h_m^H) g - \lambda g = 0, \text{ where } H_m^H f_m = h_m, \end{aligned}$$

which shows a). Part b) is trivial, since the residual of any harmonic Ritz pair is in $AK + \langle s \rangle$ and is orthogonal to AK . \square

So, to obtain harmonic Ritz pairs, we first compute f_m as the solution of an $m \times m$ linear system and then compute the eigenpairs of $H_m + f_m h_m^H$.

Note that by (3) the vector r from part b) of Lemma 1 is also the residual of the GMRES iterate with respect to the search space \mathcal{K} . Since in the situation of Lemma 1 we have

$$A\mathcal{U} \in \mathcal{U} + \langle r \rangle \tag{11}$$

for any subspace \mathcal{U} spanned by some harmonic Ritz vectors, we see that the augmented Krylov subspaces $K_j(A, r, \mathcal{U}) := \mathcal{U} + K_j(A, r)$ satisfy

$$A \cdot (\mathcal{U} + K_j(A, r)) \subset \mathcal{U} + K_{j+1}(A, r).$$

This has two major consequences. The first is that, starting from an orthonormal basis of $K_1(A, r, \mathcal{U})$, we can build nested orthonormal bases for the spaces $K_j(A, r, \mathcal{U})$ in an Arnoldi type manner. The second is that we can compute harmonic Ritz vectors with respect to $K_m(A, r, \mathcal{U})$ in the way given in Lemma 1.

This is why we are able to set up a restarted GMRES method with deflation of harmonic Ritz vectors, where each cycle consists of the following steps:

1. Extract k harmonic Ritz pairs from the “search space” $\mathcal{U}^{pr} + K_m(A, r^{pr})$ built in the previous cycle; these harmonic Ritz vectors span the current augmenting subspace \mathcal{U} .

2. With r denoting the residual of the iterate x at the beginning of the current cycle, compute nested orthonormal bases v_1, \dots, v_{k+j} for $K_j(A, r, \mathcal{U})$, $j = 1, \dots, m$, such that for $V_{k+j} = [v_1 | \dots | v_{k+j}]$ we have the Arnoldi type relation

$$AV_{k+j} = V_{k+j+1}\widehat{H}_{k+j}, \widehat{H}_{k+j} \in \mathbb{C}^{(k+j+1) \times (k+j)}. \quad (12)$$

3. Compute the current GMRES iterate and its residual using (12).

Here, as in the sequel, we use the superscript pr to denote quantities from the previous cycle; no superscript refers to the current cycle. For an efficient implementation, all three steps are coupled by the way in which we construct the orthogonal basis for the current search space $K_j(A, r, \mathcal{U})$. We now discuss the three steps in detail.

Extracting the Ritz pairs. Let $V_{\widetilde{m}}^{pr}$, $\widetilde{m} = k+m$, denote the matrix containing the orthonormal basis vectors of the search space built up in the previous cycle. By Lemma 1a), the k harmonic Ritz vectors $u_\ell, \ell = 1, \dots, k$ to augment the current Krylov subspace can be computed using $\widehat{H}_{\widetilde{m}}^{pr}$ from the Arnoldi-type relation (12) of the previous cycle. The harmonic Ritz vectors are then given as

$$u_\ell = V_{\widetilde{m}}^{pr} g_\ell, g_\ell \in \mathbb{C}^{\widetilde{m}}, \ell = 1, \dots, k,$$

which we summarize as

$$U_k := [u_1 | \dots | u_k] = V_{\widetilde{m}}^{pr} G_k, G_k = [g_1 | \dots | g_k]. \quad (13)$$

Computing the next GMRES iterate and its residual. Assume that we have already built V_{k+m+1} , the columns which represent an orthonormal basis of the current search space. We have to compute $\eta \in \mathbb{C}^{k+m}$ such that $\|b - A(x + V_{k+m}\eta)\|$ is minimal. The residual $r = b - Ax$ of the current iterate, which is in $K_1(A, r, \mathcal{U}^{pr})$ and thus in $\text{range}(V_{k+1})$, can be expressed as

$$r = b - Ax = V_{k+m+1}c, \text{ where } c = \begin{bmatrix} V_{k+1}^H r \\ \mathbf{0} \end{bmatrix}.$$

Taking (12) as granted for the moment, we have

$$b - A(x + V_{m+k}\eta) = V_{m+k+1} \left(c - \widehat{H}_{m+k}\eta \right).$$

Since the columns of V_{m+k+1} are orthonormal, minimizing $b - A(x + V_{m+k}\eta)$ is therefore equivalent to solving the small $(m+k+1) \times (m+k)$ least squares problem

$$\eta = \operatorname{argmin} \left\| c - \widehat{H}_{m+k}\eta \right\|.$$

This gives the next GMRES iterate $x^{next} = x + V_{m+k}\eta$ with residual

$$r^{next} = V_{m+k+1} \left(c - \widehat{H}_{m+k}\eta \right). \quad (14)$$

Building the orthonormal basis for the current search space. The crucial part here is to obtain an orthonormal basis for $K_1(A, r, \mathcal{U})$ without investing multiplications with A .

From the columns of $V_{\tilde{m}}^{pr}$ being orthonormal, we obtain an orthonormal basis v_1, \dots, v_k of \mathcal{U} by computing an orthonormal basis for range G_k from (13), i.e. by computing the QR-factorization $G_k = Q_k R_k$, where $Q_k \in \mathbb{C}^{\tilde{m} \times k}$ has orthonormal columns, and putting

$$V_k := [v_1 | \dots | v_k] = V_{\tilde{m}}^{pr} Q_k.$$

From (4) we have $A\mathcal{U} \subset \mathcal{U} + \langle r \rangle$, where r is the residual of the iterate at the beginning of the current cycle, $r = V_{\tilde{m}+1}^{pr} (c^{pr} - \widehat{H}_{\tilde{m}}^{pr} \eta^{pr})$ according to (14). Defining $G_{k+1} \in \mathbb{C}^{(\tilde{m}+1) \times (k+1)}$ as

$$G_{k+1} = \begin{bmatrix} G_k & c^{pr} - \widehat{H}_{\tilde{m}}^{pr} \eta^{pr} \\ 0 & \end{bmatrix}, \quad (15)$$

we thus have that $\text{range}(V_{\tilde{m}+1}^{pr} G_{k+1}) = \mathcal{U} + \langle r \rangle$. So we can extend v_1, \dots, v_k to an orthonormal basis v_1, \dots, v_{k+1} of $\mathcal{U} + \langle r \rangle$ by extending the QR-factorization $G_k = Q_k R_k$ to a QR factorization $G_{k+1} = Q_{k+1} R_{k+1}$ where

$$Q_{k+1} = \begin{bmatrix} Q_k & * \\ 0 & * \end{bmatrix} \in \mathbb{C}^{(\tilde{m}+1) \times (k+1)}.$$

This gives the relation

$$V_{k+1} = [v_1 | \dots | v_k | v_{k+1}] = V_{\tilde{m}+1}^{pr} Q_{k+1}. \quad (16)$$

On the other hand, by (11) for $j = 0$ there exists a matrix $\widehat{H}_k \in \mathbb{C}^{(k+1) \times k}$ such that

$$AV_k = V_{k+1} \widehat{H}_k. \quad (17)$$

This matrix is actually given as

$$\widehat{H}_k = Q_{k+1}^H \widehat{H}_{\tilde{m}}^{pr} Q_k, \quad (18)$$

which can be seen by comparing

$$AV_k = AV_{\tilde{m}}^{pr} Q_k = V_{\tilde{m}+1}^{pr} \widehat{H}_{\tilde{m}}^{pr} Q_k$$

which follows from the Arnoldi relation of the previous cycle and

$$AV_k = V_{k+1} \widehat{H}_k = V_{\tilde{m}+1}^{pr} Q_{k+1} \widehat{H}_k,$$

which is due to (16) and (17). We have therefore shown (12) for $j = 0$.

The Arnoldi-type relation (17) for an orthonormal basis of $\mathcal{U} + K_1(A, r)$ can now be extended to an orthonormal basis of $\mathcal{U} + K_j(A, r)$ through the standard Arnoldi orthogonalization procedure, where in step j we orthogonalize Av_{k+j}

against all previous basis vectors. As a consequence, the matrices \widehat{H}_{k+j} in the resulting Arnoldi-type relation (12) have the form

$$\widehat{H}_{k+j} = \left[\begin{array}{c|cccc} \widehat{H}_k & h_{1,k+1} & \dots & \dots & h_{1,k+j} \\ & \vdots & \vdots & \vdots & \vdots \\ & h_{k+1,k+1} & \dots & \dots & h_{k+1,k+j} \\ \hline 0_{j \times k} & h_{k+2,k+1} & \dots & \dots & h_{k+2,k+j} \\ & 0 & \ddots & \vdots & h_{k+3,k+j} \\ & \vdots & \ddots & \ddots & \vdots \\ & 0 & \dots & 0 & h_{k+j+1,k+j} \end{array} \right] \in \mathbb{C}^{(k+j+1) \times (k+j)},$$

which has upper Hessenberg structure except for the first $k+1$ rows. Herein, \widehat{H}_k is from (18); for the other entries $h_{i\ell}$ see Algorithm 4 below.

5.5. Deflating flexible GMRES

If we use a constant preconditioner M , the deflating procedure described so far can be directly extended to right preconditioned GMRES by just changing the operator from A to AM . Note that we now compute harmonic Ritz values and vectors for AM rather than A .

As was pointed out in [23, 24], going from constant to varying preconditioning, is conceptionally almost trivial: For each j the preconditioner M_j occurs only once in the matrix-vector product $z_j = M_j v_j$ of the flexible Arnoldi process. Since the vectors v_j are linearly independent—they are even mutually orthogonal—there exist a constant matrix M for which $Mv_j = M_j v_j$ for $j = 1, \dots, m$. We don't have to know M explicitly, all we need is its action on the v_j which is given through the variable preconditioning. So we can perform the deflation procedure in exactly the same way as in the case of a constant preconditioner.

Algorithmically, we have to take care of storing the preconditioned vectors $z_j = M_j v_j$ explicitly, as it is done in the flexible Arnoldi process. The resulting deflated, flexible restarted GMRES method (FGMRES-DR) is formulated as Algorithm 4.

5.6. Computing the implicit norm of the residual within GMRES-DR

As noted in section 5.2, the QR factorization of H can be updated at each step using Givens rotations. The *implicit* norm of the residual is easily computed by applying in sequence the rotations to the right hand side c of the little least squares problem and computing the absolute value of the last element of $c' = Q^H c$. The set of rotations form the unitary matrix Q^H such that $Q_{m+1}^H H_{m+1} = R_m$ with R_m triangular. The application of Q_{m+1}^H to c is easily seen as a change of basis. This new basis is composed by a set of orthonormal vectors which spans the subspace inverted by the iteration, plus a vector which is not inverted. The first m vectors of the new basis \tilde{V}_{m+1} satisfy $AZ_m = \tilde{V}_m R$. The $m+1$ element of c' is the component of b on the non-inverted,

Algorithm 4: F-GMRES-DR

Input : $A \in \mathbb{C}^{n \times n}$, $b \in \mathbb{C}^n$, restart length m , dimension of augmenting (harmonic Ritz) subspace k

- 1 choose initial guess x_0 , compute $r_0 = b - Ax^0$
 /* first cycle: flexible GMRES, no deflating subspace */
- 2 Perform m steps of the flexible Arnoldi process starting with r_0 . This gives Z_m, V_{m+1}, \hat{H}_m from (6).
- 3 Put $c = \|r_0\|e_1 \in \mathbb{C}^{m+1}$.
- 4 Compute $\eta = \operatorname{argmin}_{\xi \in \mathbb{C}^m} \|c - \hat{H}_m \xi\|$
- 5 Obtain GMRES iterate $x = x_0 + Z_m \eta$ and residual $r = r_0 - V_{m+1} \hat{H}_m \eta$
- 6 **repeat** /* all other cycles */
 - /* get augmenting (harmonic Ritz) subspace */
 - 7 Compute all m eigenpairs (λ_i, g_i) of $H_m + f_m h_m^H$ /* Lemma 1a) */
 - 8 Identify the k smallest (in modulus) harmonic Ritz values λ_i , collect the corresponding Ritz vectors g_i as columns of $G_k \in \mathbb{C}^{m \times k}$.
 /* initialization for augmented Arnoldi process */
 - 9 Build $G_{k+1} = \begin{bmatrix} G_k & c - \hat{H}_m \eta \\ 0 & \end{bmatrix}$ from (15) and compute its QR-factorization $G_{k+1} = Q_{k+1} R_{k+1}$.
 - 10 Update $\hat{H}_k := Q_{k+1}^H \hat{H}_m Q_k$, where Q_k results from Q_{k+1} by deleting its last row and column.
 - 11 Update $V_{k+1} = V_{m+1} Q_{k+1}$, $Z_k = Z_m Q_k$.
 /* Remaining part of augm. flex. Arn. */
 - 12 **for** $j = k + 1, \dots, m$ **do**
 - 13 $z_j := M_j v_j$
 - 14 $w := Az_j$
 - 15 **for** $i = 1, \dots, j$ **do**
 - 16 $h_{ij} := v_i^H w$
 - 17 $w = w - h_{ij} v_i$
 - 18 $h_{j+1,j} := \|w\|$
 - 19 $v_{j+1} := w / h_{j+1,j}$
 - /* express (old) residual in terms of new basis */
 - 20 Compute $c = \begin{bmatrix} V_{k+1}^H r \\ 0 \end{bmatrix} \in \mathbb{C}^{m+1}$ /* so $r = V_{m+1} c$ */
 /* Now obtain iterate and residual for this cycle */
 - 21 Compute $\eta = \operatorname{argmin}_{\xi \in \mathbb{C}^m} \|c - \hat{H}_m \xi\|$
 - 22 Update GMRES iterate $x = x + Z_m \eta$ and corresponding residual
 $r = r - V_{m+1} \hat{H}_m \eta$
- 23 **until** convergence

normalized vector \tilde{v}_{m+1} . In GMRES-DR the matrix H is not Hessenberg. It is however possible to compute the implicit norm of the residual by computing a QR factorization of the non-Hessenberg submatrix of H and then updating the QR factorization by Givens rotations.

5.7. Mixed precision

On most of modern CPUs, GPUs as well as on the CELL processor, single precision arithmetic is twice as fast as double precision arithmetic. The natural way of exploiting the possible acceleration provided by single precision arithmetic units, while aiming at a full double precision result when solving a linear system, is by using the technique of *iterative refinement*, see, e.g., [25].

Given two different available machine precisions, a high precision p_h and a low precision p_l , iterative refinement is based on the algorithmic principle described in Algorithm 5.

Algorithm 5: Iterative Refinement	
choose initial guess x_0	/* precision p_h */
1	compute $r_0 = b - Ax^0$ with precision p_h
2	repeat /* refinement cycles */
3	convert r_0 to precision p_l , $r \leftarrow r_0$
4	solve $Ax = r$ up to precision p_l using p_l
5	convert x to precision p_h and update $x_0 \leftarrow x_0 + x$
6	compute $r_0 = b - Ax^0$ with precision p_h ;
7	until <i>convergence</i>

In practice, where we work with IEEE single ($p_l = 2^{-24} \approx 10^{-7}$) and double ($p_h = 2^{-53} \approx 10^{-16}$) precision and where we aim at a final norm of the residual of the order of 10^{-12} , for example, two cycles are usually sufficient to obtain the desired accuracy. However, if the solver used in each refinement cycle is deflated flexible GMRES, we may encounter problems if we want to use the harmonic Ritz vectors from the last GMRES cycle of the current refinement cycle as a deflating subspace for the first GMRES cycle in the next refinement cycle: Upon convergence in low precision, the updated, low precision residual obtained via FGMRES-DR and the explicitly computed high precision residual may differ substantially. Then, including the more precise, explicitly computed high precision residual into the next deflated GMRES cycle is not possible, even when converted to low precision, because the fundamental relation (11) between the residual and the harmonic Ritz vectors, which is at the heart of the efficient use of deflation in FGMRES-DR, is lost.

To cope with this situation, we developed the following approach: Our refinement cycles correspond one-to-one to the restart cycles of FGMRES-DR, so that we have more than just the 2 or 3 usual refinement cycles. Still, of course, the overwhelming part of all computation is done in low precision arithmetic. We then recompute the residual $r = b - Ax$ of the current iterate x explicitly

in high precision and convert it back to low precision. This explicitly computed residual r is used in Algorithm 4 to obtain the right hand side c of the least squares problem that has to be solved in each cycle of FGMRES-DR (last 4 lines), and for $c - \widehat{H}_m \eta$ when obtaining the matrix G_{k+1} at the beginning of the repeat-loop. We use r also when updating the last vector of V_{k+1} . This last vector corresponds to the residual after orthonormalization against the updated V_k . Our approach can thus be regarded as an attempt to sneak in exact residual information via the right hand side of the least squares problem and one vector of the subspace while at the same time maintaining the crucial relation (11) although it will probably not hold exactly after using r to get v_{k+1} .

This approach works satisfactorily well in practice. Nevertheless, it may still happen that after some cycles the updated and the explicitly recomputed residual differ by more than a small amount. In this case, we perform a “clean” restart of the method, i.e. we start Algorithm 4 from scratch with a first cycle that does not use any deflating subspace. As a criterion for triggering the clean restart we use $\|r^e\|/\|r^i\| > t$, with r^e being the explicitly recomputed residual, r^i being the implicit residual computed as described in section 5.6, and t being a tunable threshold. From our experiments, $2 \leq t \leq 10$ is a reasonable search range.

6. Numerical results

We now report results of several numerical experiments which illustrate the impact of including deflation into the restarted flexible GMRES method. Our target systems are dynamically produced Wilson-Dirac operators with clover improvement [26]. We start with the results of a series of experiments for a relatively small $16^3 \times 32$ lattice. We used a thermalized configuration obtained from a dynamical simulation at $\beta = 5.29$ and $\kappa_{\text{sea}} = 0.135$. The value of c_{sw} in the clover improvement was 1.91; the resulting critical value for the hopping parameter κ was $\kappa_c \approx 0.13707$.

Our implementation used 4 nodes of QPACE, i.e. 32 CELL SPU cores in total. For the domain decomposition, we divided the lattice into 64 sublattices of size $4^3 \times 8$ each, so that each of the 32 cores is assigned 8 of these sublattices. We always performed 8 SAP-cycles for preconditioning, and the solves for the subdomains were done approximately via 5 steps of MR, independently of the subdomain and of the accuracy of the solution thus obtained. This choice of parameters was found to be the best by an extensive numerical study.

Figure 2 shows the results for flexible GMRES without deflation (left) and with deflation (right) for various choices of the cycle length m . Here, as everywhere, we plot the 2-norm of the residual against the number of matrix-vector multiplications invested, i.e. for every “interior” step of the restarted GMRES cycle. In these experiments the hopping parameter was set to $\kappa = 0.1368$, which is still quite far from κ_c , so that we may consider the system as relatively well conditioned. The figures illustrate the fact that a larger value of the cycle length m results in fewer iterative steps. Deflation results in fewer iterations when the

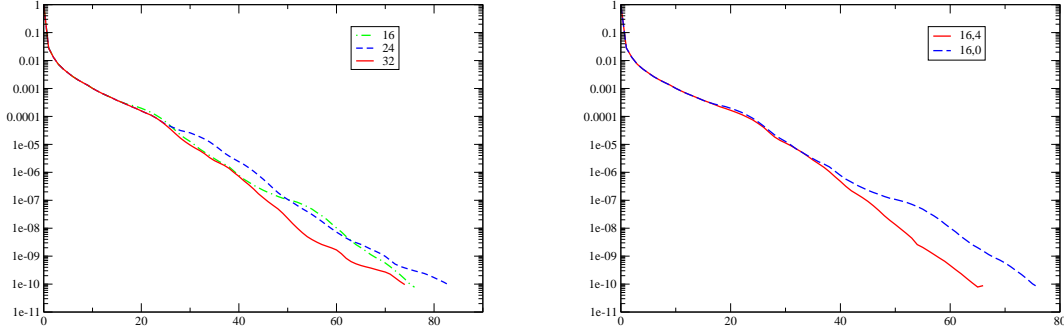


Figure 2: Results without (left) and with (right) deflation. $16^3 \times 32$ configuration, $\kappa_c \approx 0.13707$, $\kappa = 0.1368$. The legend on the left figure refers to the cycle length m ; in the right figure the pair (m, k) denotes the cycle length m and the size of the deflating subspace k .

cycle length is fixed. Note that the arithmetic cost related to the Arnoldi process grows like nm^2 due to the orthogonalizations, so larger cycle lengths m can significantly increase the computing time. This being said, Figure 2 shows that by investing just 4 vectors for deflation, we get a better performance by using a restart value $m = 16$ than by using a restart value m twice as large and no deflation.

Figure 3 presents similar experiments for a different value of the hopping parameter, $\kappa = 0.1371$, which is very close to the critical value. Now the system is much more ill conditioned, and we see that deflation starts to have an even more significant impact on the performance of the method. For example, reserving $k = 4$ vectors for deflation while using a restart value of $m = 16$ reduces the required iterations from about 450 to 115. The best deflated method uses $m = 20$ and deflates $k = 8$ vectors. It requires roughly only half as many iterations as the best non-deflated method which requires a much larger cycle length, $m = 32$. In some of the residual plots in the figure to the right we observe an increase of the residual around iterations 50-70 for some choices of the parameters. At these points we had to do a “clean” restart of FGMRES-DR, and the jump in the residual indicates that the explicitly computed residual is considerably larger than the updated one. Since a clean restart discards all subspace information acquired so far, the subsequent GMRES cycle is comparably as slow as the very first cycle of the iteration. This shows that it might be advisable to use quite small values for the subspace dimension k , since in these cases a clean restart was never necessary, thus resulting in the fastest overall convergence.

In a last experiment for this configuration we set $\kappa = 0.1374$, which is beyond the critical value. The SAP iteration by itself is divergent for this choice of κ (see Figure 5), and when used as a preconditioner for GMRES we observe

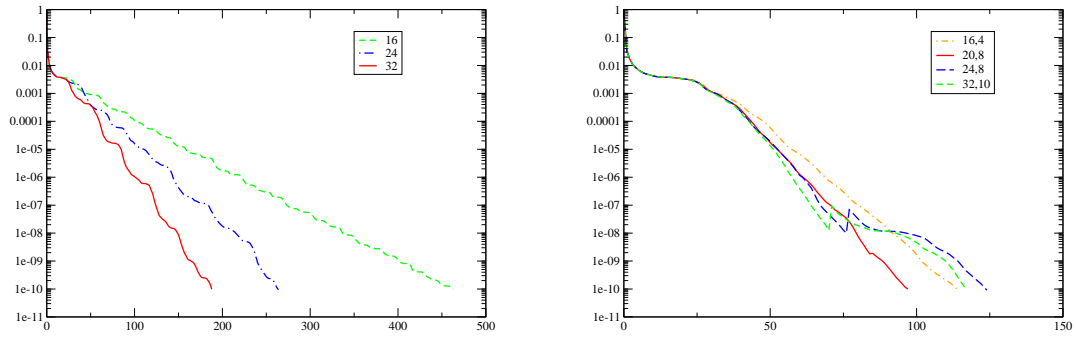


Figure 3: Results without (left) and with (right) deflation. $16^3 \times 32$ configuration, $\kappa_c \approx 0.1371$, $\kappa = 0.1371$.

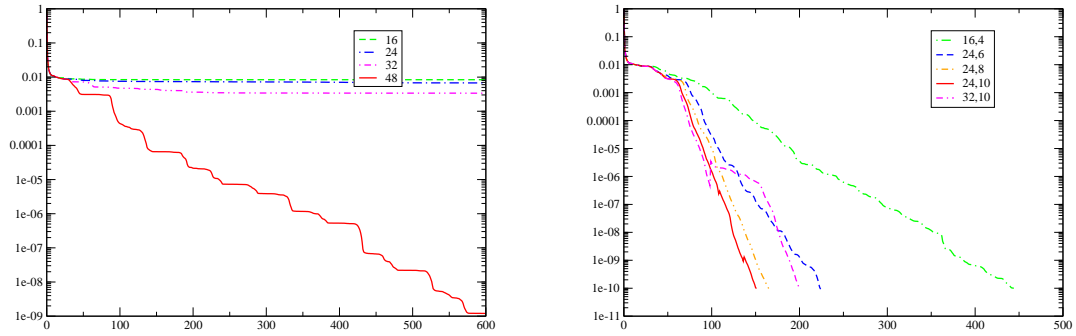


Figure 4: Results without (left) and with (right) deflation. $16^3 \times 32$ configuration, $\kappa_c \approx 0.1371$, $\kappa = 0.1374$.

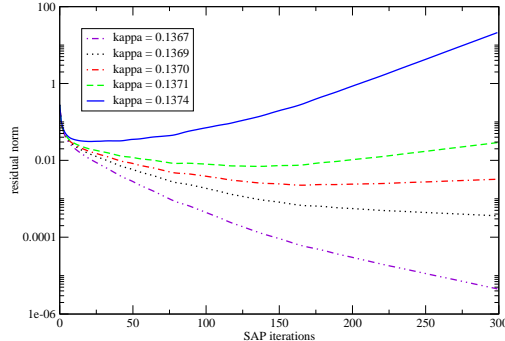


Figure 5: Convergence and divergence of the SAP method as a stand-alone solver for various values of κ

complete stagnation unless the restart value m is large enough ($m = 48$). On the other hand, deflation with a relatively small dimension of the deflating subspace results in quite fast convergence for moderate values of the restart value m . This indicates that the divergence of the pure SAP iteration is due to the fact that just a few eigenvectors belonging to the corresponding iteration matrix are troublesome and that they are effectively removed by deflation.

The second set of experiments consists of inversions of a state-of-the-art $48^3 \times 64$ $N_f = 2$ lattice configuration thermalized at $\beta = 5.40$, $\kappa_{\text{sea}} = 0.1366$ and $c_{sw} = 1.8228$, while $\kappa_c \approx 0.13670$

In this case a full 256 nodes QPACE rack was used. The machine is configured as a 3D $4 \times 8 \times 8$ torus and we used an SAP block size of $8 \times 2 \times 6 \times 6$ that fits the 8 SPU local stores nicely and satisfies all the lattice and implementation constraints. For the test inversions presented, we have chosen to use aggressive preconditioner settings that are particularly advantageous on massively parallel machines. These settings reduce the overhead of global sums and scalar products — global sums being limited by network latency and local scalar products by memory bandwidth. In the preconditioner, we therefore performed 24 SAP cycles with 6 MR iterations for the subdomains. Figure 6 shows that for $\kappa = 0.13663$, using a small dimension for the deflating subspace, i.e., $k = 3$, and keeping the restart length fixed to $m = 18$, the iteration number decreases from 183 to 159, i.e., by about 15%. For the larger value of $\kappa = 0.13666$, deflating 3 vectors halves the iteration numbers from approximately 500 down to 250. Increasing the cycle length does not substantially affect the deflated method. Figure 7 shows results for larger values of κ . For $\kappa = 0.13668$ non-deflated restarted GMRES almost stagnates even if we use a relatively large cycle length, $m = 36$. On the other hand, deflating 3 or 6 vectors cures this stagnation completely, and we can even reduce the cycle length down to 24 or even 16. A similar observation holds for $\kappa = 0.13670$.

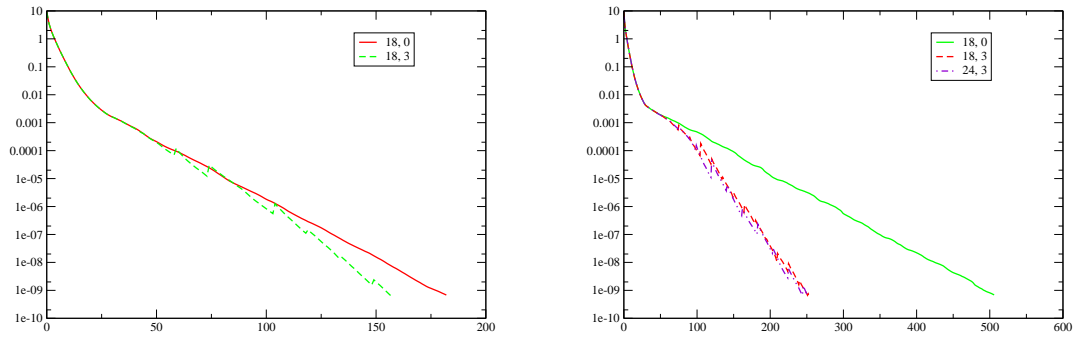


Figure 6: Results with and without deflation for different solver parameters. $48^3 \times 64$ configuration, $\kappa = 0.13663$ (left), $\kappa = 0.13666$ (right).

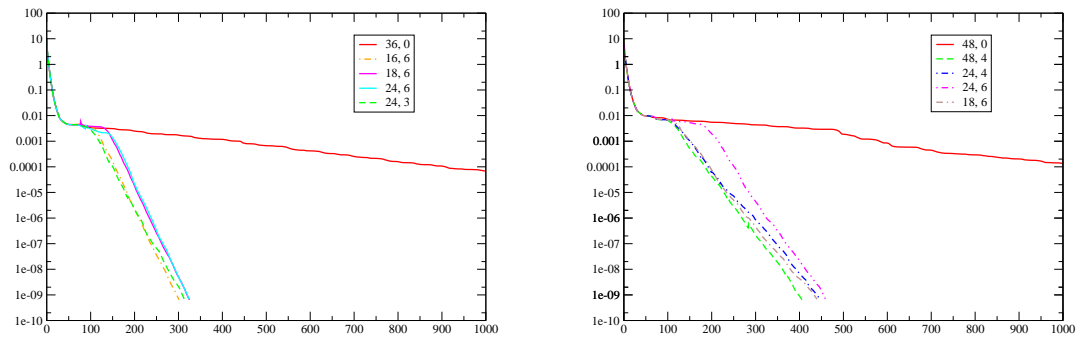


Figure 7: Results with and without deflation for different solver parameters. $48^3 \times 64$ configuration, $\kappa = 0.13668$ (left), $\kappa = 0.13670$ (right).

In Table 1 we report the timings on QPACE for all our computations on $48^3 \times 64$ configurations. These numbers show that wall clock time follows the iteration counts very closely. For the range tested, the cycle length m by itself has only a marginal influence. This is due to the fact that our code spends most of its time in the 24 SAP iterations of the preconditioner, so that the work for the Arnoldi process, which increases with m , is relatively cheap for any of our choices for m .

kappa	m	k	iterations	time (s)
0.13663	18	0	183	39.0
0.13663	18	3	159	35.9
0.13666	18	0	507	107.1
0.13666	18	3	255	56.7
0.13666	24	3	252	57.4
0.13668	36	0	3464	741.9
0.13668	16	6	302	70.4
0.13668	18	6	326	74.9
0.13668	24	6	325	73.7
0.13668	24	3	315	69.9
0.13670	48	0	2869	623.6
0.13670	48	4	409	92.3
0.13670	24	4	448	100.0
0.13670	24	6	461	106.1
0.13670	18	6	441	102.0

Table 1: Timings and iterations for the $48^3 \times 64$ configuration for the tested κ and solver settings combinations

7. Conclusions

While SAP can be implemented efficiently on current supercomputers with a deep memory hierarchy and many compute cores like QPACE, it is not necessarily a convergent iteration, particularly if the hopping parameter κ is close to the critical value. This can very severely impede the convergence of SAP preconditioned restarted flexible GMRES or restarted GCR. We have introduced *deflated* restarted flexible GMRES as a (Clover) Wilson fermion solver to cure this problem. Numerical results obtained for thermalized configurations of size up to $48^3 \times 64$ show that using a very small dimension for the deflating subspace (never more than 6) always results in satisfactory convergence speeds. An additional benefit is that deflation allows to use relatively small cycle lengths (in our experiments, $m = 16$ to 24 was always sufficient), which reduces the arithmetic cost of the Arnoldi process and, more importantly, keeps memory requirements low. The benefits of deflation are most prominent when the hopping parameter κ is close to the critical value κ_c or even larger than κ_c , so that

deflation is particularly helpful on exceptional configurations arising during an HMC simulation.

The parameter space on which the performance of SAP preconditioned deflated restarted flexible GMRES can be optimized, is very large: the size and shape of the subdomains for SAP, the method and number of iterations to use for the (approximate) subdomain solves, the number of SAP iterations per preconditioning step, the cycle length m and the deflating subspace dimension k . We could therefore not perform a complete investigation of this parameter subspace. Rather, our approach was to fix parameters related to the SAP preconditioner by efficiency considerations for our given hardware, and then find good values for k and m . Although this quite surely means that we missed the overall best method, all our results consistently indicate that it is always worth to include a (small) deflating subspace: Convergence is faster for the same cycle length, and for hard problems, convergence is enabled at all. We thus think that deflated restarted flexible GMRES should establish itself as the standard successor to flexible restarted GMRES or GCR and even more to preconditioned BiCGstab for which we observed convergence problems relatively often during simulations for large configurations.

We note that an alternative approach to the method presented here are multilevel methods like adaptive algebraic multigrid, see [27, 28] or Lüscher's deflated domain decomposition method [29] and its true multilevel variants [30]. As compared to these approaches, the deflated GMRES method does not require a set-up phase which can be relatively costly for the other methods if only one or a few right hand sides have to be solved. It has also the advantage of being quite simple to implement and to fit modern supercomputer architectures well.

Acknowledgements

We would like to thank Dirk Pleiter for his continuing support of this work and Tilo Wettig and Jaques Bloch for illuminating discussions. We are grateful to Xavier Pinel for making his FORTRAN code on FGMRES-DR available to us.

8. Bibliography

References

- [1] C. Gattringer, C. B. Lang, Quantum Chromodynamics on the Lattice, Springer, 2009.
- [2] S. Duane, A. Kennedy, B. Pendelton, D. Roweth, Hybrid Monte Carlo, Phys. Lett. B195 (1987) 216–222.
- [3] H. Baier, et al., QPACE: A QCD parallel computer based on Cell processors, PoS(Lattice 2009)001, arXiv: 0911.2174 [hep-lat].

- [4] H. Baier, et al., QPACE: Power-efficient parallel architecture based on IBM PowerXCell 8i, *Computer Science - Research and Development* 25 (2010) 149.
- [5] S. Ashby, et al., The opportunities and challenges of exascale computing, U. S. Dept. of Energy Office of Science, http://science.energy.gov/~media/ascr/ascac/pdf/reports/Exascale_subcommittee_report.pdf (2010).
- [6] M. Frigo, S. G. Johnson, The design and implementation of FFTW3, *Proceedings of the IEEE* 93 (2005) 216–231.
- [7] K. G. Wilson, Quarks and strings on a lattice, in: A. Zichichi (Ed.), *New Phenomena In Subnuclear Physics. Part A. Proceedings of the First Half of the 1975 International School of Subnuclear Physics, Erice, Sicily, July 11 - August 1, 1975*, Vol. 321 of CLNS Reports, Plenum Press, New York, 1977, p. 69.
- [8] W. Hackbusch, *Iterative Solution of Large Sparse Systems of Equations*, Springer, 1994.
- [9] B. F. Smith, P. E. Bjørstad, W. D. Gropp, *Domain Decomposition. Parallel Multilevel Methods for Elliptic Partial Differential Equations*, Cambridge University Press, 1996.
- [10] M. 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.
- [11] A. Nobile, Solving the Dirac equation on QPACE, PoS(Lattice 2010)034, arXiv:1109.4279 [hep-lat].
- [12] Y. Nakamura, A. Nobile, D. Pleiter, H. Simma, T. Streuer, T. Wettig, F. Winter, Lattice QCD applications on QPACE, arXiv: 1103.1363 [hep-lat].
- [13] Y. Saad, M. H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 7 (1986) 856–869.
- [14] O. Axelsson, A generalized conjugate gradient, least square method, *Numer. Math.* 51 (1987) 209–227.
- [15] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd edition, SIAM, Philadelphia, PA, 2003.
- [16] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, *SIAM J. Sci. Stat. Comput.* 14 (1993) 7461–469.

- [17] P. Jiránek, M. Rozložník, M. H. Gutknecht, How to make simpler GMRES and GCR more stable, *SIAM J. Matrix Anal. Appl.* 30 (4) (2008) 1483–1499.
- [18] R. B. Morgan, GMRES with deflated restarting, *SIAM J. Sci. Comput.* 24 (2002) 20–37.
- [19] A. Chapman, Y. Saad, Deflated and augmented Krylov subspace techniques, *Numer. Linear Algebra Appl.* 4 (1997) 43–66.
- [20] R. B. Morgan, A restarted GMRES method augmented with eigenvectors, *SIAM J. Matrix Anal. Appl.* 16 (1995) 1154–1171.
- [21] R. B. Morgan, Implicitly restarted GMRES and Arnoldi methods for non-symmetric systems of equations, *SIAM J. Matrix Anal. Appl.* 21 (2000) 1112–1135.
- [22] C. Paige, B. Parlett, H. V. der Vorst, Approximate solutions and eigenvalue bounds from Krylov subspaces, *Numer. Linear Algebra Appl.* 2 (1995) 115–134.
- [23] X. Pinel, A perturbed two-level preconditioner for the solution of three-dimensional heterogeneous Helmholtz problems with applications to geophysics, Ph.D. thesis, Institut National Polytechnique de Toulouse, "http://www.cerfacs.fr/algor/reports/Dissertations/TH_PA_10_55.pdf" (May 2010).
- [24] L. Giraud, S. Gratton, X. Pinel, X. Vasseur, Flexible GMRES with deflated restarting, *SIAM J. Sci. Comput.* 32 (4) (2010) 1858–1878.
- [25] J. Demmel, Y. Hida, W. Kahan, X. S. Li, S. Mukherjee, E. J. Riedy, Error bounds from extra-precise iterative refinement, *ACM Trans. Math. Softw.* 32 (2006) 325–351.
- [26] B. Sheikholeslami, R. Wohlert, Improved continuum limit lattice action for QCD with Wilson fermions, *Nucl. Phys. B* 259 (1985) 572.
- [27] J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn, C. Rebbi, Adaptive multigrid algorithm for lattice QCD, *Phys. Rev. Lett.* 100:041601, arXiv:0707.4018v1 [hep-lat].
- [28] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. A. Manteuffel, S. F. McCormick, J. C. Osborn, C. Rebbi, Adaptive multigrid algorithm for the lattice Wilson-Dirac operator, *Phys. Rev. Lett.* 105:201602, arXiv:1005.3043v2 [hep-lat].
- [29] M. Lüscher, Local coherence and deflation of the low quark modes in lattice QCD, *JHEP* 07(2007)081, arXiv:0706.2298v4 [hep-lat].
- [30] A. Frommer, K. Kahl, S. Krieg, B. Leder, M. Rottmann, Aggregation-based multilevel methods for Lattice QCD, *Proceedings of Science*, <http://pos.sissa.it>, poS(Lattice 2011)046 (2011).