

Solving DWF Dirac Equation Using Multisplitting Preconditioned Conjugate Gradient

Jiqun Tu^{*†}

Department of Physics, Columbia University, New York, NY 10027, USA

E-mail: jt2798@columbia.edu

We show that using the multisplitting algorithm as a preconditioner for conjugate gradient inversion of the domain wall fermion Dirac operator could effectively reduce the inter-node communication cost, at the expense of performing more on-node floating point operations. This method could be useful for supercomputers with far more on-node flops than inter-node communication bandwidth.

The 36th Annual International Symposium on Lattice Field Theory - LATTICE2018

22-28 July, 2018

Michigan State University, East Lansing, Michigan, USA.

^{*}Speaker.

[†]The speaker wishes to thank his advisor Robert Mawhinney for the support for this work. The author also wants to express gratitude to Guo Duo, Chulwoo Jung, Christopher Kelly and Norman Christ for the suggestions and comments. The algorithm is implemented numerically with the help from these libraries: `CPS`, `Grid`, `Qlattice` and `QUDA`. Thanks to Kate Clark and NVIDIA for the numerical support.

1. Introduction

The cost of lattice QCD simulations with dynamical fermions is dominated by the solution of the Dirac equation in both the ensemble generation phase, where configurations of gauge fields are generated, and the measurement phase, where expectation values of physical observables are measured. The Dirac matrix, which is the gauge field dependent discretization of the fermionic part of the continuous QCD action, is a large sparse linear system and inverting the corresponding Dirac equation poses tremendous numerical difficulty. For domain wall fermions(DWF) the conjugate gradient(CG) algorithm proves to be a stable algorithm to solve the Dirac equation but the convergence rate is limited by the condition number of the Dirac matrix, which is typically large in simulations with physical pion mass.

For the measurement phase various eigen-space methods, including EigCG[1] and implicitly restarted Lanczos algorithm with Chebyshev polynomial[2], have been developed successfully to speed up the inversion. Low-lying eigenvectors(eigenvectors corresponding to small eigenvalues) of the Dirac matrix are generated and the previously large condition number is effectively reduced to improve the convergence rate of CG. In this phase for one gauge field configuration typically a large number of Dirac equations with the same Dirac matrix but different right hand sides(RHS, or sources) are solved. The large number of sources amortizes the cost of eigenvector generation and the total computation time is reduced.

This is not the case for the ensemble generation phase. During a typical hybrid Monte Carlo(HMC) evolution of a gauge field as few as one Dirac equation is solved for a single Dirac matrix. This renders it not worthwhile to generate the low-lying eigenvectors for a particular Dirac matrix.

The development of supercomputers has greatly increased the number of floating point operations per second(flops) that can be performed on each processor(node). Modern lattice simulations usually divide the gauge field and pseudo-fermion fields into sub-fields that are stored and computed locally on different processors of a large parallel computer. This increases the total theoretical floating point operation capability. Inter-processor data transfer(communication), however, is needed to perform coherent operations, including the Dirac matrix multiplication. Computations locally performed on one processor require contents of the sub-fields that are stored and updated on other processors. For a specific operation if the rate of communication could not keep up with the local flops then communication becomes the bottleneck and the high flops are not utilized.

For standard CG solver with DWF one Dirac matrix multiplication is performed for each iteration. The precise requirement varies with the size of the lattice and processor grid, but roughly this requires one byte of communication for each local floating point operation. On some of the newest machines, for example the SUMMIT machine at Oak Ridge National Laboratory(ORNL), inter-processor communication speed is much less than the requirement set by their high local floating point operation capability.

In [3] a domain decomposition algorithm is proposed for Dirac equation with Wilson fermion. Local inversions are performed on two halves of the lattice iteratively. However, attempts to apply the same or similar algorithms to the inversion of the DWF Dirac equation have not been successful.

In this work we report on our investigation into a preconditioned CG solver for solving the DWF Dirac equation for the ensemble generation phase of the simulation. We find a preconditioner that decreases the number of CG iterations needed for a solution, while increasing the local

computation required per iteration, thus changing the balance of local computation to off-processor communication.

2. Method

2.1 Multisplitting Algorithm

In [4] a *multisplitting* algorithm is proposed for solving generic large linear systems distributed across a parallel computer. Compared to the domain decomposition algorithm in [3], it does not require checkerboarding. Before each iteration the boundary content of the solution field on each of the processors is communicated to its neighbors. During each iteration, the algorithm uses this communicated neighboring solution field as the Dirichlet boundary condition to perform the inversion of a local matrix on each processor. After each iteration, the updated boundary content is again communicated to prepare for the next iteration.

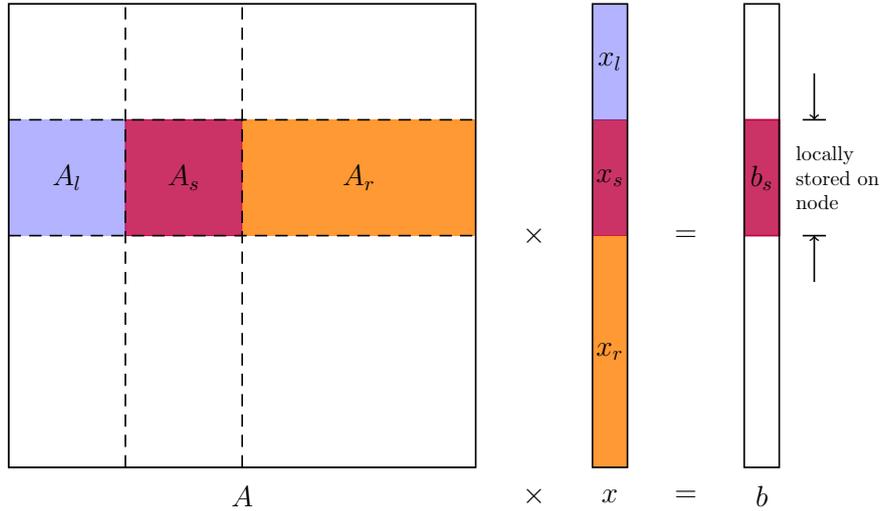


Figure 1: Decomposition of the matrix A , the solution vector x and the right-hand-side(RHS) vector b into local parts on each node.

Following [5], suppose the equation to be solved is $Ax = b$. For a *particular* processor the matrix A and vectors x and b are decomposed according to figure 1, where x_s and b_s are the part that is locally stored on this processor. On each processor the original equation turns into

$$A_s x_s + A_l x_l + A_r x_r = b_s. \tag{2.1}$$

The $A_l x_l + A_r x_r$ part involves off-processor content and is calculated before each iteration via communication. A_s is the part of the matrix that requires only the locally stored part of x on a certain processor s , i.e. x_s . Then for each iteration the algorithm solves the equation

$$A_s x_s = b_s - A_l x_l - A_r x_r \tag{2.2}$$

locally for x_s on this processor. The updated solution x_s will then be communicated to the neighboring processors. This whole procedure can be done concurrently on all nodes once the communication work to calculate $A_l x_l + A_r x_r$ is done.

2.2 Domain Wall Fermions

The domain wall fermion(DWF)[6] formulation is based on Wilson fermion and a fictitious fifth dimension. Modern numerical implementations of DWF utilize the fact that only the matrix elements that connect the *even* sites to *odd* sites and those connecting *odd* sites to *even* sites depend the gauge field. The matrix entries that connect *even* sites to *even* sites and those connect *odd* sites to *odd* sites are constant. Here the even-odd parity is defined by the 4D components of a site:

$$\text{parity} \equiv (x + y + z + t) \pmod{2}. \quad (2.3)$$

In the 4D even-odd preconditioning form the Möbius DWF Dirac equation can be written as,

$$\begin{pmatrix} M_5 & M_{eo}^4 \\ M_{oe}^4 & M_5 \end{pmatrix} \begin{pmatrix} \psi_e \\ \psi_o \end{pmatrix} = \begin{pmatrix} \phi_e \\ \phi_o \end{pmatrix}, \quad (2.4)$$

where the subscript e/o refer to even and odd sites. This is equivalent to solving the following even-odd preconditioned equation,

$$D_{PC} \psi_e = \hat{\phi}_e, \quad D_{PC} \equiv M_5 - M_{eo}^4 M_5^{-1} M_{oe}^4, \quad \hat{\phi}_e \equiv \phi_e - M_{eo}^4 M_5^{-1} \phi_o. \quad (2.5)$$

Here $M_{eo/oe}^4$ includes the Wilson hopping term $D_{x,y}^w$ that connects 4D space-time sites to their nearest neighbors,

$$M_{oe/oe}^4 = D_{x,y}^w M_\phi, \quad D_{x,y}^w \equiv \sum_{\mu} \left[(1 + \gamma_{\mu}) U_{x-\hat{\mu},\mu}^{\dagger} \delta_{x-\hat{\mu},y} + (1 - \gamma_{\mu}) U_{x,\mu}^{\dagger} \delta_{x+\hat{\mu},y} \right], \quad (2.6)$$

and M_5 and M_ϕ are constant matrices that are diagonal in the four Euclidean space-time dimensions. Details of these matrices can be found in [7].

The CG algorithm requires the matrix to be hermitian and positive definite. A common practice is to multiply both sides of (2.5) with D_{PC}^{\dagger} and solve the equation with the normal operator $D_{PC}^{\dagger} D_{PC}$ and the new RHS $D_{PC}^{\dagger} \hat{\phi}_e$ instead,

$$D_{PC}^{\dagger} D_{PC} \psi_e = D_{PC}^{\dagger} \hat{\phi}_e. \quad (2.7)$$

2.3 Dirichlet Boundary Condition on the 4-Hop Normal Operator

There are four Wilson hopping terms, one in each $M_{eo/oe}^4$, in the normal operator $D_{PC}^{\dagger} D_{PC}$,

$$D_{PC}^{\dagger} D_{PC} = [M_5 - M_{eo}^4 M_5^{-1} M_{oe}^4]^{\dagger} [M_5 - M_{eo}^4 M_5^{-1} M_{oe}^4]. \quad (2.8)$$

To apply the multisplitting algorithm to equation (2.7) Dirichlet boundary conditions are to be enforced on the normal operator $D_{PC}^{\dagger} D_{PC}$, i.e. the local part(the A_s in (2.1)) of this normal operator needs to be constructed. As the vector content is distributed across the processors according to its 4D space-time location, this local part for $D_{PC}^{\dagger} D_{PC}$ includes *snake* terms that hop out of the

boundary and hop back in as the various components in (2.8) are evaluated. Figure 2 illustrates this and gives some examples of the snake terms. These terms are truncated if Dirichlet boundary conditions are enforced on each of the four $M_{eo/oe}^4$ hopping terms sequentially. Our simulation results show that the inclusion of these snake terms is crucial to the convergence.

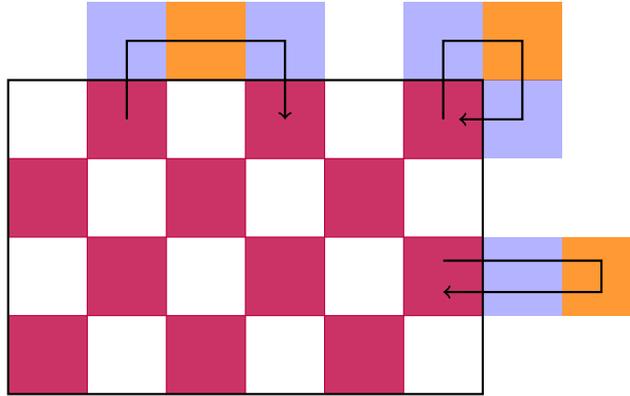


Figure 2: The normal operator $D_{PC}^\dagger D_{PC}$ has as many as 4 Wilson hopping terms. Enforcing Dirichlet boundary condition on it requires the inclusion of the *snake* terms, e.g. the black arrows.

2.4 Multisplitting Algorithm as a Preconditioner of CG

In [3] to achieve faster convergence the domain decomposition algorithm is eventually used as a preconditioner of GCR. In this work we use the multisplitting algorithm as a preconditioner of CG.

Pseudocode for a generic preconditioned CG is shown below, where we are solving $Ax = b$ and M is the preconditioning matrix. The preconditioning step is marked with blue background. The overall convergence rate of preconditioned CG is estimated by the condition number of AM^{-1} . If the condition number of AM^{-1} is smaller than that of the original matrix A , faster convergence rate is achieved.

Now for this preconditioning step we use the multisplitting algorithm to solve for z_{k+1} in

$$Az_{k+1} = r_{k+1}. \quad (2.9)$$

To avoid inter-processor communication, a zero initial guess ($x_l = x_r = 0$) is used in (2.2) and only the first iteration is performed. With r_{k+1} as the RHS and z_{k+1} the solution,

$$A_s x_s = b_s - A_l x_l - A_r x_r \rightarrow A_s z_{k+1,s} = r_{k+1,s}. \quad (2.10)$$

This is equivalent to using the local part of the matrix A , A_s , on each processor as the preconditioner M in the preconditioned CG,

$$M = \bigoplus_s A_s, \quad s = \text{node index}. \quad (2.11)$$

The local nature of A_s makes it possible to perform the preconditioning step concurrently on all the processors without communication. We refer to this as multisplitting preconditioned CG(MSPCG).

Algorithm 1 Preconditioned Conjugate Gradient $Ax = b$

$$r_0 = b - Ax_0$$

$$z_0 = M^{-1}r_0$$

$$p_0 = z_0$$

$$k = 0$$

while have not converged **do**

$$\alpha_k = \langle r_k, z_k \rangle / \langle p_k, Ap_k \rangle$$

$$x_{k+1} = x_k + \alpha_k p_k$$

$$r_{k+1} = r_k - \alpha_k Ap_k$$

$$z_{k+1} = M^{-1}r_{k+1}$$

$$\beta_k = \langle z_{k+1}, r_{k+1} \rangle / \langle z_k, r_k \rangle$$

$$p_{k+1} = z_{k+1} + \beta_k p_k$$

$$k = k + 1$$

end while

3. Results

The multisplitting preconditioned CG is applied to solve Dirac equations on three 2+1 flavor lattice ensembles generated with Möbius domain wall fermions, all with physical input quark masses. Standard CG is used to perform the inversion in the preconditioning step. Instead of adopting a precision based stopping condition, a fixed number of CG iterations, which will be referred as *inner iterations*, are performed for these preconditioning solves. The iterations performed in the overall preconditioned CG will be referred as *outer iterations*. In table 1 the numbers of outer iterations needed for the preconditioned CG to converge are reported on the different lattice ensembles, together with the stopping condition for the outer CG(precision) and the processor grid size used. The numbers of iterations to reach the same precision with standard CG are also included for comparison, where the inner iteration number is marked with *plain*.

Typically on these ensembles with 6 inner iterations the preconditioned CG reduces the outer iteration count by a factor of 3. More inner iterations reduce the outer iteration count more but the reduction saturates as the inner iteration count increases: with large number of inner iterations the inner CG solves the preconditioning inversion completely and no further numerical benefit can be exploited.

4. Conclusion

Our results show the MSPCG reduces the number of outer iterations needed to solve the DWF Dirac equation, reducing the inter-processor communication at the expense of performing more local inner iterations. We observe that executing a fixed number of inner CG iterations for the preconditioning inversion, instead of using a precision based stopping condition, does not jeopardize the convergence of the outer CG. This is true even when as few as 3 inner iterations are performed. As a consequence the inner iteration count is a parameter that can be tuned to achieve maximum speed up in the trade-off between inter-processor communication and local computation.

lattice size	$a^{-1}[\text{GeV}]$	precision	processor grid size	inner iterations	outer iterations
$32^3 \times 64$	1.37	10^{-8}	—	plain	13594
			$2^3 \times 4$	3	9106
			$2^3 \times 4$	4	6020
			$2^3 \times 4$	6	5126
$64^3 \times 128$	2.36	10^{-10}	—	plain	18092
			$4^3 \times 8$	6	6008
			$4^3 \times 8$	12	5083
			$4^3 \times 8$	18	4948
$80^2 \times 96 \times 192$	3.00	10^{-10}	—	plain	16783
			$4^2 \times 8^2$	6	5719

Table 1: Number of outer iterations need to converge the multisplitting preconditioned CG for the lattice ensembles tested in this work. *Inner iterations* refers to the fixed number of CG iterations performed for the preconditioning inversion. Rows marked with *plain* indicate the iteration count for the same standard CG to converge.

We note that while the multisplitting algorithm can split the general matrix A in a variety of ways, the splitting presented here, used as a preconditioner in CG, makes it equivalent to the additive Schwarz algorithm. (The additive Schwarz algorithm has been used for the Dirac equation inversion for the fermions[8, 9].) We use the name MSPCG, as it is through the process of applying the multisplitting algorithm to the DWF Dirac equation that we realize the necessity of including the snake terms in the local matrix.

References

- [1] A. Stathopoulos and K. Orginos, SIAM J. Sci. Comput. **32**, 439 (2010).
- [2] Y. Saad, SIAM J. Numer. Anal. **17**, 687 (1980).
- [3] M. Lüscher, Comput. Phys. Commun. **156**, 209 (2004), [arXiv:hep-lat/0310048].
- [4] D. P. O’Leary and R. E. White, SIAM J. Algebr. Discret. Methods **6**, 630 (1985).
- [5] F. Jezequel, R. Couturier, and C. Denis, J. Supercomput. **59**, 1517 (2012).
- [6] K. Jansen, Phys. Rep. **273**, 1 (1996), [arXiv:hep-lat/9410018].
- [7] R. C. Brower and H. Neff, 1 (2014), [arXiv:1206.5214].
- [8] Y. Osaki and K.-i. Ishikawa, 1 (2010), [arXiv:1011.3318].
- [9] R. Babich et al., (2011), [arXiv:1109.2935].