

Initial guesses for multi-shift solvers

James Osborn*

Argonne National Lab

E-mail: osborn@alcf.anl.gov

I will present a method for providing initial guesses to a linear solver for systems with multiple shifts. This can also be extended to the case of multiple sources each with a different shift.

The XXVI International Symposium on Lattice Field Theory

July 14 - 19, 2008

Williamsburg, Virginia, USA

*Speaker.

1. Introduction

A fundamental part of lattice QCD calculations is the solution of a discretized Dirac equation

$$(D+m)x = b \quad (1.1)$$

for some source field b . Here D is the Dirac matrix (for some choice of discretization), m is the quark mass (times the identity) and x is the desired solution. This is typically solved with methods such as Conjugate Gradient (CG), that find a solution among the Krylov space $\{b, (D+m)b, (D+m)^2b, \dots\}$.

Often it is necessary to solve this equation for several masses against the same source. This can be done efficiently with a class of Krylov methods that solve for all shifts at the same time [1]. These methods are made possible because the Krylov spaces for different shifts still span the same linear space, and the solutions can all be obtained in a single pass of the algorithm. The number of iterations required for the solution of all equations is then just the number needed for the worst conditioned equation (lightest mass).

Since solving the Dirac equation can make up a large part of lattice QCD calculations, it becomes very important to find ways to reduce the time needed to get a solution. Often one can use some prior knowledge about similar systems to the one being solved to obtain initial guesses which, in the case of a single shift, can easily be used to reduce the number of iterations needed. The prior knowledge can be from previous solutions at a lower precision, projection of low eigenmodes (or approximate ones), or solutions of similar systems with small changes in either the source or the matrix (such as in the chronological inverter [2]).

Unfortunately for systems with multiple shifts the use of the prior information is not as simple. This is due to the residuals obtained from the guesses not being the same in general. Here we will present a method that can use this information to produce initial guesses with a common right hand side so that standard multi-shift Krylov methods can still be used. The problem of initial guesses is related to the more general problem of solving systems with multiple shifts each with a different source, which we will also provide an algorithm for. While the method of initial guesses does provide an improvement in some cases, a straightforward implementation may in other cases produce initial residuals that are too large to be useful. We will show examples that demonstrate this breakdown and discuss some possible methods to alleviate it.

2. Multiple shift solvers and initial guesses

Here we are interested in solving the system of N linear equations

$$(A + \sigma_i) x_i = b \quad (1 \leq i \leq N) \quad (2.1)$$

where A is matrix and σ_i are shifts of a constant times the identity. As mentioned in the introduction, these equations can be solved simultaneously by using multi-shift Krylov methods that exploit their common Krylov space. These multi-shift methods form the solutions from the common Krylov space $\{b, Ab, A^2b, \dots\}$. If one wanted to make use of some initial guesses, y_i , for the solutions to reduce the number of iterations, the typical thing to do is construct

$$r_i = b - (A + \sigma_i) y_i \quad (2.2)$$

and then solve

$$(A + \sigma_i) z_i = r_i . \quad (2.3)$$

The solutions would be then given by $x_i = y_i + z_i$. However, in general, the new right hand sides, r_i , are not collinear, and hence they won't share the same Krylov space, preventing the use of multi-shift Krylov methods which would solve them all at the same time.

There is a simple form for the guesses that will make the r_i the same. Take

$$y_i = \left[\prod_{j \neq i} (A + \sigma_j) \right] w \quad (2.4)$$

for any vector w , then one can easily see that the new right hand sides are all equal to

$$r_i = b - \left[\prod_j (A + \sigma_j) \right] w . \quad (2.5)$$

The system can then be solved with standard multi-shift Krylov methods.

The problem now is just to find the best choice for w . Consider first the case of $N = 2$ shifts. Given approximate solutions

$$\begin{aligned} v_1 &\approx (A + \sigma_1)^{-1} b \\ v_2 &\approx (A + \sigma_2)^{-1} b \end{aligned} \quad (2.6)$$

with corresponding residuals

$$\begin{aligned} R_1 &= b - (A + \sigma_1) v_1 \\ R_2 &= b - (A + \sigma_2) v_2 \end{aligned} \quad (2.7)$$

then a good choice for w could be

$$w = (v_1 - v_2) / (\sigma_2 - \sigma_1) \approx [(A + \sigma_1)(A + \sigma_2)]^{-1} b \quad (2.8)$$

giving

$$r_1 = r_2 = [(A + \sigma_2)R_1 - (A + \sigma_1)R_2] / (\sigma_2 - \sigma_1) . \quad (2.9)$$

Note that if v_1 and v_2 were exact solutions then starting residual would be zero.

For general N the corresponding choice for w would be

$$w = \sum_i c_i v_i \quad (2.10)$$

with

$$c_i = \prod_{j \neq i} \frac{1}{\sigma_j - \sigma_i} . \quad (2.11)$$

Note that the coefficients c_i can become large as one goes to more shifts with smaller differences. As we will see later, this can lead to a breakdown of the algorithm if care is not taken to keep the common residual (2.5) from growing too large.

3. Multiple shifts with multiple sources

It turns out that the problem of initial guesses is a special case of the more general case of multiple shifts each with a different source, which can also be solved. The two-source two-shift method was worked out in [3]. Consider the system

$$\begin{aligned}(A + \sigma_1) x_1 &= b_1 \\ (A + \sigma_2) x_2 &= b_2 .\end{aligned}\tag{3.1}$$

Now choose guesses y_k such that the residuals are equal

$$b_1 - (A + \sigma_1) y_1 = b_2 - (A + \sigma_2) y_2 .\tag{3.2}$$

By equating powers of A we find

$$y_1 = y_2 = (b_2 - b_1) / (\sigma_2 - \sigma_1)\tag{3.3}$$

which gives a common starting right hand side of

$$b_i - (A + \sigma_i) y_i = [(A + \sigma_2) b_1 - (A + \sigma_1) b_2] / (\sigma_2 - \sigma_1)\tag{3.4}$$

which is just (2.9) with the b_i replaced by R_i .

To extend this to arbitrary N we need to find a set of y_i that give a common residual r

$$b_i - (A + \sigma_i) y_i = r \quad (1 \leq i \leq N) .\tag{3.5}$$

This can be solved by setting

$$y_i = \sum_{j=0}^{N-2} A^j s_{i,j} \equiv p_i(A)\tag{3.6}$$

then equating powers of A and solving for the vectors $s_{i,j}$. One can also solve this by considering the polynomials $q_i(A) = (A + \sigma_i)p_i(A)$ at the special cases of $A = -\sigma_k$ where the residual $r = b_k$. This gives the N equations (for fixed i)

$$q_i(-\sigma_k) = b_i - b_k .\tag{3.7}$$

Since $q_i(A)$ is a polynomial of order $N-1$ in A the system is uniquely determined. The polynomial satisfying these equations is

$$q_i(A) = \sum_k \left[\prod_{j \neq k} \frac{A + \sigma_j}{\sigma_j - \sigma_k} \right] (b_i - b_k)\tag{3.8}$$

which gives

$$y_i = \sum_{k \neq i} \left[\prod_{j \neq i,k} \frac{A + \sigma_j}{\sigma_j - \sigma_k} \right] \frac{b_i - b_k}{\sigma_i - \sigma_k} .\tag{3.9}$$

m_1	d	no guess	$N = 1$	$N = 2$	$N = 3$	$N = 4$	$N = 5$	$N = 6$
0.010	2	683	334 (-3.0)	445 (4.6)	488 (11.0)	509 (16.3)	518 (20.4)	523 (23.3)
0.010	$\sqrt{2}$	683	334 (-3.0)	489 (5.4)	575 (13.1)	635 (20.3)	677 (26.8)	985 (32.7)
0.005	2	1365	666 (-3.0)	892 (5.8)	978 (13.5)	1018 (19.9)	1039 (25.3)	1214 (29.4)
0.005	$\sqrt{2}$	1365	666 (-3.0)	977 (6.6)	1151 (15.6)	1272 (23.9)	1761 (31.7)	3066 (38.8)
0.002	2	3417	1668 (-3.0)	2228 (7.4)	2445 (16.6)	2554 (24.7)	3965 (31.6)	6723 (37.3)
0.002	$\sqrt{2}$	3417	1668 (-3.0)	2445 (8.2)	2885 (18.8)	3246 (28.7)	7210 (38.0)	10000 (46.8)
0.001	2	6830	3353 (-3.0)	4464 (8.6)	4903 (19.0)	5328 (28.3)	10000 (36.4)	10000 (43.3)
0.001	$\sqrt{2}$	6830	3353 (-3.0)	4886 (9.4)	5764 (21.1)	8949 (32.3)	10000 (42.8)	10000 (52.7)

Figure 1: Iterations (initial $\log_{10}|r_i|^2$) for guesses formed from approximate solutions.

4. Initial tests

To demonstrate the strengths and weaknesses of this method, we have performed some simple tests. For all tests we are using an even-odd preconditioned “asqtad” staggered Dirac matrix so that we are solving the Hermitian positive definite system

$$(m_k^2 - D_{eo}D_{oe}) x_k = b \quad (4.1)$$

where D_{eo} and D_{oe} are the even-odd and odd-even blocks of the Dirac matrix and the shift is now the square of the masses. The source vector b is taken to be a point source. For simplicity in all tests we used a random gauge field with an average plaquette value of around 0.39 (normalized to 1). The masses are set to be geometrically spaced, $m_k = m_1 d^{k-1}$, with $m_1 \in \{0.01, 0.005, 0.002, 0.001\}$ and $d \in \{2, \sqrt{2}\}$. The final stopping criterion for the residual is $|r|^2 < 10^{-6}$ (the source is normalized to 1). While this is a fairly relaxed criterion, it was chosen to keep the number of iterations from growing too large at the lightest mass. There are many factors that can effect the performance of this algorithm, so these tests only serve to show the qualitative behavior as more and lighter masses are used. All work was done in double precision.

The first tests are with initial guesses made from approximate solutions on a 32^4 lattice. The approximate solutions were obtained from running multi-shift CG until $|r|^2 < 10^{-3}$. These solutions were then used to generate the initial guesses from (2.10) and (2.4). This example is done purely for testing purposes since the residuals $|r_i|^2$ for $i \geq 3$ had already converged to the final precision. Also since the guesses came from another multi-shift CG, their residuals could have already been collinear, which we could have taken advantage of as discussed later.

In figure 1 we show the results for the approximate solutions. The “no guess” column gives the number of iterations necessary when starting with zero guess. The N value is the number of equations (shifts) solved simultaneously. In those columns are the number of iterations needed before the accumulated residual from the CG reached the stopping criterion (with a maximum of 10,000 iterations) along with the value of $\log_{10}(|r_i|^2)$ for the initial common residual (2.5) used for the new right hand side. After the CG stopped, the true residual was calculated for all shifts. The numbers in orange and red indicate that the true residuals had actually not converged, with orange for $10^{-6} < |r|^2 < 10^{-5}$ and red for $10^{-5} < |r|^2$.

m_1	d	no guess	$N = 1$	$N = 2$	$N = 3$	$N = 4$	$N = 5$	$N = 6$
0.010	2	665	626 (0.1)	650 (6.7)	677 (13.0)	693 (18.3)	705 (22.4)	708 (25.3)
0.010	$\sqrt{2}$	665	626 (0.1)	657 (7.1)	700 (14.5)	740 (21.5)	772 (28.0)	1078 (33.9)
0.005	2	1328	1136 (0.3)	1185 (8.1)	1241 (15.6)	1280 (22.1)	1303 (27.4)	1605 (31.5)
0.005	$\sqrt{2}$	1328	1136 (0.3)	1200 (8.5)	1288 (17.1)	1369 (25.3)	1946 (32.9)	3263 (40.0)
0.002	2	3309	2157 (0.4)	2268 (9.5)	2402 (18.5)	2539 (26.5)	4776 (33.4)	7549 (39.1)
0.002	$\sqrt{2}$	3309	2157 (0.4)	2291 (9.7)	2479 (19.6)	3339 (29.2)	7274 (38.4)	10000 (47.0)
0.001	2	6563	3627 (0.5)	3800 (10.1)	4028 (20.1)	5969 (29.3)	10000 (37.4)	10000 (44.3)
0.001	$\sqrt{2}$	6563	3627 (0.5)	3805 (10.3)	4099 (21.3)	8618 (32.2)	10000 (42.6)	10000 (52.5)

Figure 2: Iterations (initial $\log_{10} |r_i|^2$) for guesses formed from approximate eigenmodes.

For all cases that converged, the number of iterations was less with the guess than without. Remarkably even for a starting residual of $|r_i|^2 \approx 10^{21}$ the residual could be reduced to the final precision in fewer iterations. However as one moves toward more or smaller masses, the initial residual grows very large until it is no longer possible to reduce it all the way back to the final goal in double precision.

In the second set of tests the guesses were obtained by projection of approximate eigenmodes of the preconditioned Dirac matrix. Here the lattice size was 16^4 . The low modes were obtained simply by repeated inversions on random vectors with occasional Rayleigh-Ritz diagonalization. The final vectors were still far from the lowest eigenmodes since the smallest approximate eigenvalue (Ritz value) was still at least 4 times larger than the lowest true eigenvalue. This was done to give a more difficult test of the algorithm since with exact eigenvalues the deflation can be done exactly and the starting residuals are automatically equal.

In figure 2 we show the results for the approximate eigenmodes with the same conventions as the previous table. Again we see the same pattern of improved convergence up to the point that the initial residual becomes too large to reduce in double precision. The only exception is at the heaviest mass where the low mode projection is no longer effective anyway. Clearly the method is providing good guesses for the low modes of the system. The main difficulty then is keeping the initial residual under control so that the solver can converge. Next we will discuss some possible strategies for this.

5. Variations

For the problem of choosing initial guesses with the same right hand side, there are several possible strategies for choosing the vector w used in (2.10). One possibility is to globally optimize for w from $r = b - (A + \sigma_1) \dots (A + \sigma_N)w$ among some given search space of vectors. This can be done by either minimizing the norm of the residual or by projecting out the search space from the residual. An alternative is to individually optimize each equation separately, $R_k = b - (A + \sigma_k)v_k$, then apply the multi-source multi-shift algorithm to get the initial guesses. In practice the latter seems to give better guesses, though global minimizations control the residual better. By

interpolating between the two one can then find the best compromise that still leads to a solution. Of course at some point this will be no better than an initial guess of zero.

Another variation is given by the observation that the starting residuals don't have to be equal, but merely collinear. Thus we can add arbitrary scale factors to the b_i in (3.9). This is especially useful if the trial guesses are obtained from another run using, e.g., CG. Here the residuals would be collinear in exact precision, and in finite precision may be close, but not exact. Restarting with the appropriate scale factors could give a large improvement in this case.

6. Conclusions

We have presented a method for solving systems with multiple sources each with a different shift. The main motivation was to provide initial guesses to multi-shift solvers, though it could be useful in other contexts as well. When used for initial guesses we found that even though the initial residuals may be large, the convergence is still typically faster as long as convergence can still be reached. The method breaks down at some point when going to more and/or smaller shifts. This can be remedied at the expense of using a worse initial guess, which may still reduce the number of iterations overall in some cases. A better solution to this problem may require projecting out the high eigenmodes of the residual while preserving the low modes of the guesses.

References

- [1] R. W. Freund, *Solution of shifted linear systems by quasi-minimal residual iterations*, in *Numerical Linear Algebra*, L. Reichel, A. Ruttan and R.S. Varga (eds.), 1993; U. Glässner, S. Güsken, T. Lippert, G. Ritzenhöfer, K. Schilling and A. Frommer, *How to compute Green's functions for entire mass trajectories within Krylov solvers*, hep-lat/9605008; A. Boriçi, *Krylov subspace methods in lattice QCD*, SCSC report TR-96-27; B. Jegerlehner, *Krylov space solvers for shifted linear systems*, hep-lat/9612014.
- [2] R.C. Brower, T. Ivanenko, A.R. Levi and K.N. Orginos, *Chronological inversion method for the Dirac matrix in hybrid Monte Carlo*, *Nucl. Phys. B* **484** (1997) 353.
- [3] M.A. Clark, *The rational hybrid Monte Carlo algorithm*, Ph.D. Thesis, University of Edinburgh (2006).