TWO-GRID DEFLATED KRYLOV METHODS FOR LINEAR EQUATIONS

RONALD B. MORGAN[†] AND ZHAO YANG[‡]

Abstract. An approach is given for solving large linear systems that combines Krylov methods with use of two different grid levels. Eigenvectors are computed on the coarse grid and used to deflate eigenvalues on the fine grid. GMRES-type methods are first used on both the coarse and fine grids. Then another approach is given that has a restarted BiCGStab (or IDR) method on the fine grid. While BiCGStab is generally considered to be a non-restarted method, it works well in this context with deflating and restarting. Tests show this new approach can be very efficient for difficult linear equations problems.

AMS subject classifications. 65F10, 15A06

1. Introduction. We look at solving large systems of linear equations Ax = b that come from discretization of partial differential equations. There are a variety of iterative methods for solving these problems. In particular, multigrid methods [15, 6, 7, 8, 32] are extremely effective under certain conditions. However, there are many situations for which standard multigrid methods do not work. Here we give a different approach to using the power of different grid levels for such situations. This will be done by combining with Krylov subspace methods.

Convergence of iterative methods is generally affected by the conditioning of the linear equations, or more specifically by the presence of small eigenvalues. One reason that multigrid methods are effective is that eigenvectors corresponding to small eigenvalues generally have similar shape on different grids levels, and multigrid is able to handle the small eigenvalues on coarse grids where both the linear equations problem is better conditioned and the computation is cheaper due to the smaller matrix.

Krylov subspace methods are more robust than multigrid in the sense that they can be applied to problems for which multigrid fails. However, they can converge slowly. Work has been done on dealing with the detrimental presence of small eigenvalues for Krylov methods. Restarted methods such as GMRES are particularly sensitive to the presence of small eigenvalues. Deflated GMRES methods [21, 19, 10, 3, 9, 16, 23, 27, 18, 17] compute approximate eigenvectors and use them to remove or deflate the effect of small eigenvalues. In particular, we will use the method GMRES-DR [23] which both solves linear equations and simultaneously computes eigenvectors.

Once approximate eigenvectors have already been computed, one deflation technique is to use them to build a preconditioner for the linear equations [19, 3, 9, 16, 27, 29]. However, such a preconditioner can become expensive if there are many eigenvalues being deflated; each eigenvector is applied at every iteration. Deflation can be done with less expense outside of the Krylov iteration [25, 31]. We will use such a method called GMRES-Proj [25] which projects over approximate eigenvectors in between cycles of regular GMRES.

^{*}The first author was supported by the National Science Foundation under grant DMS-1418677. [†]Department of Mathematics, Baylor University, Waco, TX 76798-7328 (Ronald_Morgan@baylor.edu).

[‡]Department of Mathematics, Baylor University, Waco, TX 76798-7328 (zhao_yang@alumni.baylor.edu).

R. B. MORGAN and Z. YANG

The computation of approximate eigenvectors can be expensive for difficult problems. We look at finding them by computing eigenvectors on a coarse grid and moving them to the fine grid. GMRES-DR is used on the coarse grid to both find approximate eigenvectors and generate an approximate linear equations solution. Then GMRES-Proj takes what has been found by GMRES-DR and solves the fine grid linear equations.

See Elman, Ernst and O'Leary [11] for another way of combining multigrid and GMRES. Closer to the current work is Sifuentes' thesis [28] which has two-grid deflation, but with Arnoldi on the coarse grid and with the more expensive approach of building a deflating preconditioner for GMRES on the fine grid. Instead of using computed approximate eigenvectors, Erlangga and Nabben [13] deflate using vectors from the interpolation operator that maps from coarse to fine grid.

For fairly sparse matrices, the orthogonalization expense in GMRES-Proj can be a major expense. We present an approach that substitutes either BiCGStab or IDR in place of GMRES. Here we restart BiCGStab and IDR even though they are normally non-restarted methods. This may seem risky, because of the inconsistent convergence of these nonsymmetric Lanczos methods. However, in our experiments, restarted BiCGStab and IDR converge reliably.

Section 2 reviews some of the previous methods that will be used. Section 3 presents the method Two-grid Deflated GMRES. Two-grid Deflated BiCGStab/IDR is then given in Section 4. Further examples are in Section 5 with a Helmholtz example that includes multigrid preconditioning and with an example of improving eigenvectors on the fine grid.

2. Review. Here we very quickly describe methods that will be used in the rest of the paper.

2.1. GMRES-DR. The GMRES with Deflated Restarting (GMRES-DR) [23] method uses Krylov subspaces to both solve linear equations and compute the eigenpairs with smallest eigenvalues. Once eigenvectors converge far enough, their presence in the subspace can essentially remove or deflate the effect of the small eigenvalues on the linear equations. Usually only moderate accuracy is needed before the approximate eigenvectors have a beneficial effect.

GMRES-DR(m,k) saves k approximate eigenvectors at the restart and builds out to a subspace of dimension m. Specifically, for one restarted cycle it uses subspace

$$Span\{y_1,\ldots,y_k,r_0,Ar_0,A^2r_0,\ldots,A^{m-k-1}r_0\},\$$

where y_i 's are harmonic Ritz vectors from the previous cycle and r_0 is the residual vector at the start of the cycle. This augmented subspace is actually a Krylov subspace itself, and it contains Krylov subspaces with each y_i as starting vector. This makes the eigenvectors generally converge along with the linear equations.

GMRES-DR converges faster than restarted GMRES for difficult problems with small eigenvalues. It also often converges faster than BiCGStab in terms of matrixvector products, but has greater orthogonalization costs per matrix-vector product.

2.2. GMRES-Proj. There are situations where approximate eigenvectors are available at the beginning of the solution of linear equations. For example, if there are multiple right-hand sides, then eigenvectors could have been computed during solution of earlier right-hand sides [27, 25, 31]. The method GMRES-Proj [25] uses these approximate eigenvectors to deflate the corresponding eigenvalues while solving linear equations. GMRES(m)-Proj(k) assumes that k approximate eigenvectors have

been previously computed and alternates projection over these vectors with cycles of GMRES(m).

GMRES(m)-Proj(k)

- **0.** Let k be the number of approximate eigenvectors that are available. Choose m, the dimension of subspaces generated by restarted GMRES.
- 1. Alternate between A) and B) until convergence:
 - **A)** Apply Galerkin projection over the subspace spanned by the k approximate eigenvectors.
 - **B)** Apply one cycle of GMRES(m).

The algorithm for the projection step is given next.

Galerkin projection over a set of approximate eigenvectors

- **0.** Let the current system of linear equations be $A(x x_0) = r_0$.
- 1. Let V be an n by k orthonormal matrix whos columns span the set of approximate eigenvectors.
- **2.** Form $H = V^T A V$ and $c = V^T r_0$.
- **3.** Solve Hd = c, and let $\hat{x} = Vd$.
- 4. New approximate solution is $x_p = x_0 + \hat{x}$, and new residual is $r = r_0 A\hat{x} = r_0 AVd$.

MinRes projection can be used instead of Galerkin. It is the same as Galerkin, except $H = (AV)^T AV$ and $c = (AV)^T r_0$. Often Galerkin is best and can even be much better, however we have seen situations for which MinRes works better. All experiments in this paper use Galerkin.

2.3. Two-grid Arnoldi. A two-grid method for computing eigenvalues and eigenvectors is given in [26]. Eigenvectors are computed on a coarse grid with a standard Arnoldi method, are moved to the fine grid (with spline interpolation), then are improved on the fine grid with Arnoldi-E [22], a method that can accept initial approximations.

3. Two-grid Deflated GMRES. Our new methods generate approximate eigenvectors from the coarse grid and use them to deflate eigenvalues on the fine grid. In this section, we give a version using GMRES methods. GMRES-DR is applied on the coarse grid. This generates eigenvectors, and solves the coarse grid linear equations. This linear equations solution is moved to the fine grid with spline interpolation and used there as the initial guess. The eigenvectors are similarly moved to the fine grid and, if necessary, improved on the fine grid. GMRES-Proj is applied on the fine grid using these approximate eigenvectors. This can have much faster convergence than restarted GMRES. Compared to running GMRES-DR on the fine grid, it is cheaper to implement and can deflate eigenvalues from the beginning.

Two-grid Deflated GMRES

- **0.** Choose m and k for the coarse grid. Pick nev, the number of eigenpairs that are required to converge to an eigenvalues tolerance, say rtolev. For the fine grid, pick the linear equations residual tolerance, rtol. Also choose m3 for GMRES on the fine grid.
- 1. Apply GMRES-DR(m,k) on the coarse grid. Move the approximate eigenvectors to the fine grid (we use spline interpolation). Move the solution of the coarse

grid linear equations problem to the fine grid and use it as an initial guess for the fine grid problem.

- (If needed:) Improve approximate eigenvectors on the fine grid using Arnoldi-E (see Two-grid Arnoldi method in [26]).
- **3.** Apply GMRES(m3)-Proj(k) on the fine grid.

For the first examples there is no need for the second phase, but it is used in Subsection 5.3.

Example 1. We consider a system of linear equations from finite difference discretization of the 2-D convection-diffusion equation $-e^{5xy}(u_{xx}+u_{yy})+40u_x+40u_y = c \sin x \cos x e^{xy}$ on the unit square with zero boundary conditions and c chosen to make the right-hand side be norm one. The discretization size is $h = \frac{1}{512}$, leading to a matrix of dimension $n = 511^2 = 262, 121$. The coarse grid discretization size is $h = \frac{1}{64}$ giving a matrix of dimension $63^2 = 3969$.

The first phase on the coarse grid uses GMRES-DR(150,100) and runs until 80 eigenpairs have converged to a level of residual norm below 10^{-8} . These residual norms are computed only at the end of cycles. The eigenvectors are moved to the fine grid and are accurate enough there to be effective in deflating eigenvalues (after the Rayleigh-Ritz procedure is applied to all 100 vectors that are moved from the coarse grid, the smallest 80 Ritz pairs have residual norms at or below 1.4×10^{-3}). So as mentioned, the second phase is not needed. The third phase of solving the fine grid linear equations is stopped when the relative residual norm drops below tolerance of $rtol = 10^{-10}$. This is also checked only at the end of each GMRES cycle, but of course could be easily monitored during the GMRES runs.

The top of Figure 3.1 has convergence of both linear equations and eigenvalues on the coarse grid. The linear equations converge before any of the eigenvalues, and it actually takes much longer for all 80 eigenpairs to become accurate. The linear equations converge in 19 cycles of GMRES-DR(150,100) which use 1050 matrix-vector products (150 for the first cycle and 50 each for the next 18). The eigenvalues take 107 cycles or 5450 matrix-vector products.

The bottom of the figure shows convergence of the linear equations on both grids versus the number of fine-grid-equivalent matrix-vector products. The coarse grid matrix is about 64 times smaller than the fine grid matrix, so we scale the number of matrix-vector products by a factor of 64 to get the fine-grid-equivalents. The coarse grid linear equations then converge so rapidly that the convergence curve is barely noticeable on the left of the graph. Then there is a small gap before the curve for fine grid convergence starts. This gap is for both the convergence of the eigenpairs on the coarse grid and the matrix-vector products needed to form the projection matrix Hfor the fine grid. Three different values of m3 are used for GMRES(m3)-Proj(100): 50, 100 and 200. While the fastest convergence is with GMRES(200), the least expensive is for m3 = 100. We define the approximate cost as cost = 5 * mvp + vops, where mvpis the number of matrix-vector products, the 5 comes from the approximate number of non-zeros per row, and vops is the number of length-*n* vector operations such as dot-products and daxpy's. The costs for the entire process, including the first phase on the coarse grid, is $cost = 1.10 \times 10^6$ for m3 = 50, 1.06×10^6 for m3 = 100 and finally $1.30 * 10^6$ with $m_3 = 200$. The orthogonalization cost is high for GMRES(200) and this is the biggest expensive for a sparse matrix such as this one.

Next, we consider different sizes for the coarse grid. Smaller coarse grid matrices mean less work is needed to find the eigenpairs, however they may not be as accurate



FIG. 3.1. This figure has Two-grid Deflated GMRES for the convection-diffusion example. Fine grid matrix size is n = 261, 121 and coarse grid matrix size is n = 3969. The top portion has GMRES-DR(150,100) on the coarse grid; linear equations convergence is shown with circles at the end of each cycle and convergence of 80 eigenpairs is shown with lines. The bottom portion has both the coarse grid linear equations solution on the very left, scaled by 64 to correspond to fine-grid matrix-vector products. Then this is followed by GMRES(m3)-Proj(100) on the fine grid, with m3 = 50, 100, 200.

for the fine grid work. Table 3.1 has coarse grids from size $65,025 = 255^2$ down to $49 = 7^2$. All of the tests use the same type coarse grid computation as before and use GMRES(100)-Proj(100) on the fine grid. The results show that for this matrix, the fine grid convergence is fairly robust with respect to the coarse grid size. The number of fine grid cycles increases by less than half as the coarse grid goes from size 65,025 down to 255. For the 255 case, the accuracy of 80 eigenpairs on the fine grid is $5.7 * 10^{-3}$ or better, and that is enough to be fairly effective. The convergence is 10 times faster than with no deflation which is on the last row of the table. The larger coarse grid matrices do give more accurate eigenvectors on the fine grid, for instance with residual norms $4.4 * 10^{-5}$ and below for size 65,025. However, this greater accuracy is not really needed and only speeds up the fine grid convergence from 94 to 85 cycles compared to size 3969.

Now we discuss other methods. Multigrid is much faster than Krylov when convection is low. However, here with convection terms $40u_x + 40u_y$, standard multigrid methods do not converge. Next we compare to other Krylov methods. The top part of Figure 3.2 has convergence in terms of matrix-vector products. GMRES-DR(150,100) is run just on the fine grid, as are IDR and BiCGStab. The new Two-grid Deflated GMRES method is given with m3 = 100 and it converges faster than fine grid GMRES-DR(150,100), because approximate eigenvectors from the coarse grid allow deflation of small eigenvalues from the beginning. The new method also converges in many less matrix-vector products than the non-restarted methods BiCGStab [33] and IDR [30, 34]. However, for this quite sparse matrix, the new method is not necessarily better overall than the non-restarted methods. The bottom of Figure 3.2 has convergence methods.

	Г	ABLE 3	.1	
Effect	of	Coarse	Grid	Size

coarse grid	coarse grid	accuracy of 80	fine grid	$\cos t$
matrix size cycles		eigenpairs on f.g.	cycles	(millions)
$65,025 = 255^2$	1023	$4.4 * 10^{-5}$	85	8.02
$16,129 = 127^2$	270	$3.1 * 10^{-4}$	86	1.40
$3969 = 63^2$	107	$1.4 * 10^{-3}$	94	1.06
$961 = 31^2$	43	$5.8 * 10^{-3}$	106	1.15
$225 = 15^2$	12	$5.7 * 10^{-3}$	120	1.30
$49 = 7^2$	1	$(9.2 * 10^{-3} \text{ for } 40)$	177	1.91
no coarse grid	-	-	1255	13.2



FIG. 3.2. Two-grid deflated GMRES (with GMRES-DR(150,100) on the coarse grid and GMRES(100)-Proj(100) on the fine grid) compared to GMRES-DR(150,100), BiCGStab and IDR(4). The top half is plotted against matrix-vector products and the bottom half against the cost in terms of vector op's.

gence as a function of the approximate cost and BiCGStab is much faster because of using only seven *vops* per mvp. IDR(4) uses 12.8 *vops* per mvp, but is not effective for this matrix.

We have seen in this example that Two-grid Deflated GMRES can be better than BiCGStab in terms of matrix-vector products. So it can be the better method if the matrix is not very sparse or if an expensive preconditioner is used. Then the number of iterations is the main factor instead of the GMRES orthogonalization expense.

However, for sparse problems, BiCGStab has an advantage. This motivates the next method. For the fine-grid portion of the two-grid approach, GMRES is replaced by BiCGStab. This is designed to have both a low number of matrix-vector products and low costs for vector operations.

4. Two-grid Deflated BiCGStab.

4.1. The algorithm. We wish to use approxiate eigenvectors from the coarse grid to deflate eigenvalues from BiCGStab or IDR on the fine grid. In [1, 24, 2] a deflated BiCGStab method is given. A single projection is applied before running BiCGStab using both right and left eigenvectors. Here a single projection will not be effective, because our eigenvectors are not accurate on the fine grid. So we implement BiCGStab/IDR as a restarted method with projections at each restart. Not only does this allow us to use less accurate eigenvectors, but also it does not require left eigenvectors. We use the same Galerkin projection as for Two-grid Deflated GMRES.

We now give the implementation of this restarted, deflated BiCGStab/IDR method It replaces phase 3 in the Two-grid Deflated GMRES algorithm given earlier and thus is part of a Two-grid Deflated BiCGStab/IDR algorithm. The new BiCGStab(ncyc)-Proj(k) is similar to GMRES(m3)-Proj(k), but ncyc is the total number of BiCGStab cycles, not the length of the cycles (so there are ncyc - 1 restarts). As before, the k gives the number of approximate eigenvectors that are being projected over. The stopping test for the algorithm uses the minimum of two different quantities, but this means the maximum convergence between the two criteria. First, we require each cycle to converge to a fraction of the remaining distance to final convergence, in terms of orders of magnitude. This fraction is determined by the number of remaining cycles to equalize the amount expected of each cycle. Second, for cycle *icyc* we want to reach a point that is at least *icyc/ncyc* of the way toward convergence (again in orders of magnitude). The first criteria usually asks for convergence to a further point and so is the one enforced, but in the case where the residual norm jumps up during the projection, the second criteria is needed (see Example 6).

BiCGStab(ncyc)-Proj(k) and IDR(ncyc)-Proj(k)

- **0.** Assume k approximate eigenvectors are provided. Let rtol be the specified relative residual tolerance for the linear equations solution. Choose ncyc, the requested number of cycles of BiCGStab/IDR. Let r_0 be the initial residual for the fine grid iteration.
- **1.** For icyc = 1 : ncyc
 - a) Apply Galerkin projection over approximate eigenvectors.
 - **b)** Let ||r|| is the current residual. Set the relative residual tolerance for this cycle, *rticyc*, to be the minimum (the further convergence point) of $(rtol * ||r_0||/||r||)^{(\frac{1}{ncyc-icyc+1})}$ and $(||r_0||/||r||) * (rtol)^{\frac{icyc}{ncyc}}$.
 - c) Run BiCGStab or IDR with relative residual tolerance of *rticyc*.
 - d) Break out of the loop if ||r|| is already below *rtol*.

For our tests, the Matlab BiCGStab program is called. For IDR, we use the program described by van Gijzen and Sonneveld in [34] and available in MATLAB code from the authors. The default version IDR(4) is called.

Example 2. We return to the same convection-diffusion problem as in Example 1 that has n = 262, 121 and coarse grid matrix of size 3969. Figure 4.1 has plots of BiCGStab(ncyc)-Proj(100) with ncyc = 5, 10 and 20. The more frequent restarts, and thus more deflations of the eigenvalues, allow ncyc = 20 to converge faster. It is surprising that in spite of the very jagged behavior of the residual norms of BiCGStab, the overall convergence with ncyc = 20 is quite consistent. With even more frequent restarts, the convergence is similar to ncyc = 20. These tests are not shown on the figure, because they mostly overlie the ncyc = 20 curve. However, see Table 4.1 for the number of matrix-vector products needed with some other values of ncyc. Note that with large values of ncyc, there may be convergence before all cycles are used



FIG. 4.1. Two-grid Deflated BiCGStab with GMRES-DR(150,100) on the coarse grid and BiCGStab(ncyc)-Proj(100) on the fine grid. The number of cycles for the restarted BiCGStab phase is ncyc = 5, 10 and 20. The color changes with each new cycle.

and the break out of the loop in part d) of the algorithm is activated. This is shown in parentheses in the table. For example, only 168 runs are needed when 200 are specified. We also tried restarting BiCGStab after 272 matrix-vector products (the average length for the ncyc = 20 test), and the convergence was a little slower than the ncyc = 20 method, using 5865 matrix-vector products instead of 5421. Probably the ncyc = 20 test has an advantage because it often restarts after an iteration where the residual norm comes down significantly.

Figure 4.2 has a comparison of different methods. The top portion has convergence with respect to matrix-vector products. Two-grid Deflated BiCGStab uses ncyc = 20 and has projection over 100 approximate eigenvectors. These vectors come from the same run of GMRES-DR(150,100) on the coarse grid as is used for Two-grid GMRES. This deflated BiCGStab converges faster than Two-grid Deflated GMRES. Both deflated methods are much better than regular BiCGStab. Also given is a Two-grid Deflated BiCGStab result for solving a second system of linear equations with the same matrix but a different right-hand side (this right-hand side is generated randomly). It converges similarly to the first right-hand side, but does not need the coarse-grid work or forming of the projection matrix. The bottom of Figure 4.2 has a plot of convergence versus cost. Two-grid BiCGStab converges faster than regular BiCGStab in spite of getting a late start due to the coarse grid phase. The second right-hand side does not have this cost and so has less than half the expense.

4.2. Effectiveness of restarted BiCGStab. Figure 4.1 and Table 4.1 show something remarkable for deflated, restarted BiCGStab. For $ncyc \ge 20$, the results are fairly close to being invariant of the number of cycles. This is in spite of much smaller subspaces being used for the larger values of ncyc. Also, the convergence is at a very consistent pace considering the usual erratic convergence of BiCGStab. Deflated GMRES in Figure 3.1 is very sensitive to the subspace sizes, with m3 = 50



FIG. 4.2. Two-grid Deflated BiCGStab, with GMRES-DR(150,100) on the coarse grid and BiCGStab(20)-Proj(100) on the fine grid, compared to other methods. The Two-grid Deflated GM-RES has GMRES-DR(150,100) on the coarse grid and GMRES(100)-Proj(100) on the fine grid. A test of deflated BiCGStab with a second right-hand side is also shown.

TABLE 4.1

Comparing different values of ncyc for BiCGStab(ncyc)-Proj(100). The number of matrix-vector products is given.

ncyc	5	10	15	5 20 30 50		100	150	200	
							(use 99)	(140)	(168)
mvp's	9638	7606	6390	5421	5357	5278	5262	5591	5541

converging three times slower than for m3 = 200. Deflated BiCGStab with 20 cycles uses an average of 271 matrix-vector products per cycle, while with ncyc = 100, an average of 53 per cycle converges slightly faster.

We next investigate the reason for this effectiveness of restarted BiCGStab. First, it is not due to the deflation. To show this, the next example does not use deflation and has a similar phenomenon for restarted BiCGStab.

Example 3. The matrix is the same as in the previous examples, except the size is $127^2 = 16, 129$ and there is not a coarse grid matrix. We use a random right-hand side to try to make the example more general, though generally there is not much effect from the right-hand side. We run the restarted BiCGStab with no deflation between cycles. Table 4.2 shows tests with different numbers of cycles. Surprisingly, even for large numbers of cycles, the results are similar and often better than for regular non-restarted BiCGStab. This behavior is reminiscent of Tortoise and the Hare behavior for GMRES [12] where smaller subspaces give faster convergence. This is less surprising for BiCGStab, because of the lack of the minimal residual property. However, in one way the behavior here is more extreme than Tortoise and the Hare for GMRES, because the convergence with frequent restarts can be even better than with no restarting. For example, with *ncyc* specified to be 400, the method ends TABLE 4.2

Restarted BiCGStab with different numbers of cycles (no deflation is used). The number of matrix-vector products is given.

ncyc	1 (no	5	10	25	50	75	150	250	400
	restart)					(use 74)	(118)	(171)	(198)
mvp's	3969	3755	4371	3968	3386	3578	4049	3764	3452

up using 198 cycles, 3452 matrix-vector products and the average number of matrix-vector products per cycle of $17\frac{1}{2}$. This compares to 3969 matrix-vector products with non-restarted BiCGStab. Also for GMRES(18), 28,838 matrix-vector products are required.

It has been proposed that GMRES should use changing cycle lengths [4]. We suggest that variable cycle lengths is the main reason that restarted BiCGStab is so effective even with small subspace sizes. Table 4.3 has results for restarted BiCGStab with average subspace dimensions of approximately 50, 35 and 18 (the lowest we could get was about 18, because when a high number of cycles is specified, the algorithm finishes in many fewer cycles than requested). The number of matrix-vector products is similar for each of the three average subspace sizes (these three numbers come from Table 4.2 with ncyc requested to be 75, 150 and 400). The next column has GMRES(m) for m = 50,35 and 18. As mentioned, the number of matrix-vector products is very large for small m. The third result is for GMRES restarted exactly as BiCGStab is (so with the same number of matrix-vector products for each corresponding cycle, though it does not necessarily need all of the cycles). Now GMRES uses even less matrix-vector products than BiCGStab. This is a remarkable improvement over regular restarted GMRES. Regular GMRES(m) can get stuck in a pattern [5, 35] and changing the sizes of the subspaces can break the pattern. The final column of the table has a different way of implementing variable restarting for GMRES. A maximum cycle length is specified and the length of each individual cycle is randomly chosen between 1 and the max length (for example, max cycle length of 36 is used in the test because it gives an average of 18). This approach performs much better than the fixed cycle length approach, but is not as good as using the BiCGStab lengths. We conclude from these tests that our BiCGStab method restarts in a surprisingly effective way. It is important for the new Two-grid Deflated BiCGStab method that frequent restarts still give an effective method, because if only a few restarts were used, there may not be frequent enough deflations, as seen in Figure 4.1 with the ncyc = 5 and 10 cases.

Tables 4.2 and 4.3 make it seem like the sizes of the subspaces used for BiCGStab and GMRES, restarted as for BiCGStab, do not have much effect for this example. This would be surprising, since it is well known that large subspaces can be an advantage for Krylov methods. And indeed, even though GMRES does well with these small subspaces, it is significantly better without restarting, using 1383 matrix-vector products. Since non-restarted BiCGStab has 3969 matrix-vector products, we see that here GMRES is better able to take advantage of a large non-restarted Krylov subspace.

5. Further Experiments.

5.1. Helmholtz matrix.

average	restarted	GMRES(m),	GMRES restarted	GMRES with
cycle length	BiCGStab	m fixed	as was BiCGStab	random restarts
50	3578	6250	3266	3498
35	4049	11,760	3244	3733
18	3452	$28,\!838$	3411	5013

TABLE 4.3 Matrix-vector products for restarting BiCGStab and GMRES.

Example 4. We next test the Helmholtz matrix from finite difference discretization of the differential equation $-u_{xx} - u_{yy} - 100^2 u = f$ on the unit square with zero boundary conditions. The discretization size is again $h = \frac{1}{512}$ and again n = 262, 121. The *rtol* is 10^{-10} . The right-hand side is generated random normal. The coarse grid discretization size is $h = \frac{1}{128}$, giving a matrix of dimension $127^2 = 16, 129$. This problem is difficult because of the significantly indefinite spectrum. Standard multigrid methods do not converge, however multigrid can be used as a preconditioner; see the next example. Figure 5.1 has results for Krylov methods plotted against matrix-vector products. We have found that IDR performs better than BiCGStab for indefinite matrices and only it is shown on the figure, but here even IDR does not converge. GMRES-DR(200,150) converges eventually (though not quite to the requested *rtol*), however it is expensive due to orthogonalization costs. Our Two-grid Deflated GMRES method first uses GMRES-DR(200,150) on the coarse grid until 120 eigenvalues converge to residual norm below 10^{-8} . This takes 285 cycles. Then on the fine grid it has GMRES(100)-Proj(150). This method converges faster than GMRES-DR(200,150) and is much less expensive, because on the fine grid GMRES(100) has less orthogonalization. However, Two-grid Deflated IDR is the best method. It uses IDR(20)-Proj(150) on the fine grid. It converges faster in terms of matrix-vector products and also is much less expensive per matrix-vector product.

We did try only deflating 100 eigenvalues as in Examples 1 and 2. However, deflated IDR converges more than twice as slow. For this difficult indefinite problem, many eigenvalues are needed for effective deflation.

Helmholtz problems are fairly complicated. For example, if the wave number is increased and the fine grid uses the same discretization, then the coarse grid may need to be finer than in Example 4 in order to have good enough eigenvector approximations. Our goal here is merely to show potential for the new approach; much more work is needed on Helmholtz problems for a thorough study.

Example 5. For this example, the problem is the same as in the previous one, but we now use multigrid preconditioning. Since multigrid does not converge for this matrix, the preconditioner has solution of linear equations from the positive shifted Laplacian with operator $-u_{xx} - u_{yy} + 100^2 u$ [20]. With this positive shift, multigrid easily converges and thus can precondition the Helmholtz matrix (the negatively shifted Laplacian). This preconditioning is used on both the fine grid and the coarse grid and makes the problem easier to solve; see Figure 5.2. The problem is still indefinite and while BiCGStab still does not converge, now IDR does converge. GMRES-DR(150,100) converges in less than half the number of matrix+preconditioner applications than IDR, but has more orthogonalization expense. The two-grid deflated methods use GMRES-DR(150,100) on the coarse grid, stopping when 80 eigenpairs have converged to residual norm 10^{-8} . Then on the fine grid, deflated GMRES uses GMRES(100)-Proj(100) and deflated IDR uses IDR(5)-Proj(100) (the IDR converges



FIG. 5.1. Matrix is from the simple Helmholtz equation with $\kappa = 100$. Two-grid deflated GMRES uses GMRES-DR(200,150) on the coarse grid and GMRES(100)-Proj(150) on the fine grid. Two-grid deflated IDR uses GMRES-DR(200,150) on the coarse grid and IDR(20)-Proj(150) on the fine grid; diamonds show the residual norm at the end of each of the 20 cycles. Also compared are IDR and GMRES-DR(200,150).

almost 25% slower with ncyc = 20 instead of 5). Both deflated methods converge faster than the other methods, but the deflated IDR has much less orthogonalization.

In this experiment, the positively shifted Laplacian linear equations are solved accurately. However, one can relax that and apply the multigrid only to partial convergence. The methods still work, but the results vary some.

We also tried complex shifts for the multigrid preconditioner [14]. There is faster convergence, but greater cost per iteration due to subspaces becoming complex. For a complex Helmholtz problem, this would not be a disadvantage and should be a subject for further study.

5.2. Biharmonic Matrix. We next consider matrices from discretizing a biharmonic differential equation. Matrices from this differential equation quickly become very ill-conditioned as the discretization size gets small. The biharmonic examples demonstrate first that residual norms can jump up during the projection over approximate eigenvectors. Then the second example has faster convergence if approximate eigenvectors are improved on the fine grid.

Example 6. The partial differential equation is $-u_{xxxx} - u_{yyyy} + 40u_{xxx} = f$ on the unit square with zero boundary conditions. The discretization has h = 1/256, giving a matrix of size n = 65,025. The right-hand side is chosen random normal. Due to the ill-conditioning, all Biharmonic tests have residual tolerance for the linear equations of only 10^{-8} . BiCGStab and IDR do not converge and GMRES-DR is slow and expensive. We only give results for deflated BiCGStab. The top half of Figure 5.3 shows results with three choices of m and k for GMRES-DR(m,k) on the coarse grid. GMRES-DR(100,50) finds 40 eigenpairs to residual norms below 10^{-8} , GMRES-DR(150,100) gets 80 to that tolerance and GMRES(200,150) stops when 120



FIG. 5.2. Matrix is from the simple Helmholtz equation with $\kappa = 100$. Now multigrid preconditioning is used. Two-grid deflated GMRES uses GMRES-DR(150,100) on the coarse grid and GMRES(100)-Proj(100) on the fine grid. Two-grid deflated IDR uses GMRES-DR(150,100) on the coarse grid and IDR(5)-Proj(100) on the fine grid; diamonds show the residual norm at the end of each of the 5 cycles. These methods are compared to IDR and GMRES-DR(150,100).

have converged. Deflated BiCGStab is used on the fine grid with ncyc = 20 cycles. The convergence is plotted against cost for both coarse and fine grid phases. Here each matrix-vector product is counted as 13 vector operations (the number of non-zeros in most rows). The cost for the coarse grid phase is greater when more eigenvalues are computed, but then the convergence is faster on the fine grid. Using 100 approximate eigenvectors is best overall for this example. If there were multiple right-hand sides, then 150 would be clearly better for subsequent right-hand sides because of the faster fine grid convergence.

The lower half of Figure 5.3 has a portion of the fine grid restarted BiCGStab with k = 150, plotted against matrix-vector products. The residual norm jumps up by a significant amount during each projection over the approximate eigenvectors. For instance, it increases from $2.5 * 10^{-6}$ to $1.2 * 10^{-5}$ in between cycles 14 and 15 (at matrix-vector product 3415). As a result, the second of the two convergence criteria in part 1. b) of the algorithm is activated. Figure 5.4 shows eigencomponents of residual vectors for a smaller version of this problem. The matrix is size n = 961 and the coarse grid matrix is size 49. Here 10 eigenvectors are computed accurately on the coarse grid and moved to the fine grid. The top of the figure shows all 961 eigencomponents during part of a deflated BiCGStab run on the fine grid. The (red) circles are from the residual vector after the projection at the start of the fifth cycle. Then the (black) squares are after BiCGStab has been applied. Finally, the (blue) dots are after the next projection at the start of the sixth cycle, and they mostly overlie the circles but are a little better on average. Most of the components increase dramatically with the projection. Then, fortunately they are reduced by the Krylov iteration. The lower part of the figure shows that some of the components corresponding to the small eigenvalues are reduced by the projection. This reduction is important



FIG. 5.3. Matrix is from a biharmonic differential equation. The top half has deflated BiCGStab, with convergence shown on both the coarse and fine grids. The coarse grid has GMRES-DR(m,k) for three choices of m and k. The fine grid has BiCGStab(50)-Proj(k) for the three values of k. The bottom half has a view of part of the fine grid convergence with k = 150, showing the jump in residual norm with each projection.

because they are for the most part not reduced by the BiCGStab. This essential reduction of the small eigencomponents by the projection makes up for the increase of the other components, because overall this deflated method does better than regular non-restarted BiCGStab.

5.3. Preconditioned Biharmonic. We continue the biharmonic example, but now use an incomplete LU factorization. This example also considers improving the eigenvectors on the fine grid (phase 2 of the Two-grid Deflated GMRES algorithm).

Example 7. We use incomplete factorization preconditioning for the biharmonic differential equation from the previous example. This is done with Matlab's "ilu" command with no fill, after adding 0.5 to all diagonal elements of the matrix (this is needed to make the preconditioning effective). We choose a smaller discretization size than the previous example, because the preconditioning allows us to solve a harder problem. The matrix has $n = 511^2 = 261,121$ and the course grid is size $127^2 = 16,129$. An ILU preconditioner is generated for both the fine and course grids. As in the previous example, rtol is set at 10^{-8} .

Table 5.1 has results of a few tests. The second row of the table ("cycles phase 1") has the number of GMRES-DR(150,100) cycles on the coarse grid. The third row has the cycles of Arnoldi-E on the fine grid used to improve the approximate eigenvectors (phase 2. of the Two-grid Deflated GMRES algorithm). The fourth row has the number of cycles on the fine grid; this happens to be exactly 200 for GMRES-Proj and is specified to be 50 for restarted BiCGStab. The fifth row of the table has the number of applications of matrix-vector product plus preconditioner on the fine grid. Then the next row adds to this the coarse grid mvp's+preconditionings, scaled by 16 to give the fine-grid-equivalent total. The last row has the approximate cost which counts 26 for each matrix-vector product plus preconditioner (13 for the



FIG. 5.4. Small matrix from a biharmonic differential equation. The top portion of the figure has eigencomponents of the residual vector before applying BiCGStab in the fifth cycle (red circles), then after the BiCGStab (black squares), and then after the projection over the approximate eigenvectors (blue dots). The bottom portion has only the first 20 eigencomponents.

mvp and 13 for preconditioning) and then adds the vector ops. The coarse grid costs are included, but again scaled by 16. The first result in the table (the second column) is for regular, non-restarted BiCGStab. For convergence, this takes 34,907 applications of a matrix-vector product plus a preconditioning. The next column has Two-grid Deflated GMRES, with GMRES-DR(150,100) on the coarse grid and GMRES(50)-Proj(100) on the fine grid. This is much better than BiCGStab in terms of the numbers of applications of matrix-vector product plus preconditioner, using the fine-grid-equivalent if 10,203 of them. However, the cost estimate in the last row of the table is only a little better (854 thousand compared to 1152 thousand) due to the orthogonalization in the 200 cycles of GMRES(50). If GMRES(100)-Proj(100) is used instead, the number of cycles on the fine grid reduces to 73, and thus the fine-grid equivalent matrix-vector product plus preconditioners goes down to 7576. However the cost goes up to 988 thousands. The next column has Deflated BiCGStab. The same GMRES-DR run is used on the coarse grid and then the fine grid has BiCGStab(50)-Proj(100), with 50 cycles of BiCGStab and projections over the 100 approximate eigenvectors in between. The cost is significantly reduced compared to the other methods. We note that with ncyc = 20, there are 12,168 mvp+prec's. For this example, more restarts are needed, perhaps because the approximate eigenvectors are not as good and so need to be deflated more often. The next to last column has Arnoldi-E(150,100) applied to the approximate eigenvectors for 10 cycles with the 10 smallest approximate eigenvectors used as starting vectors. Then the last column has the Arnoldi-E applied as in [26] until the first 80 approximate eigenvectors have residual norms below 10^{-3} (before this improvement, the 80 have residual norms from 3.3×10^{-4} to 2.4×10^{-2}). This improvement takes 41 cycles. Both of these raise the cost but reduce the BiCGStab iterations. For cases with multiple right-hand sides, this eigenvector improvement could be worthwhile. Improving eigenvectors was generally

Table 5.1

Results for biharmonic matrix preconditioned by ILU factorization. The first method is regular non-restarted BiCGStab. Second is Two-grid Deflated GMRES with GMRES-DR(150,100) on the coarse grid and GMRES(50)-Proj(100) on the fine grid. The last three columns have the same coarse grid work and then restarted BiCGStab(50)-Proj(100) on the fine grid. The first of these three tests has no improvement of the approximate eigenvectors on the fine grid, and the last two have 10 and 41 cycles of improvement respectively.

fine grid method	BiCGSt	G-Proj	Defl. Bi.	Defl. Bi.	Defl. Bi.
cycles phase 1	-	31	31	31	31
cycles phase 2	-	0	0	10	41
cycles phase 3	1	200	50	50	50
f. g. mvp + prec	34,907	10,000	7877	6074	3069
total f. g. equiv. mvp+	34,907	10,203	8080	6827	5372
cost (in thousands)	1152	854	322	801	2369

not beneficial in the earlier examples. Probably the best way to decide if eigenvectors should be improved is by experimenting.

6. Conclusion. We have proposed a two-grid method that finds approximate eigenvectors with the coarse grid and uses them to deflate eigenvalues for linear equations on the fine grid. This includes deflation for BiCGStab and IDR using only approximate right eigenvectors and novel use of restarting for these normally non-restarted nonsymmetric Lanczos methods. This two-grid deflation is a very efficient way to deflate eigenvalues, because the difficult work of finding approximate eigenvectors is done for an easier problem on the coarse grid. This is particularly useful for multiple right-hand side problems, because the coarse grid work only needs to be done once and then can be applied for all right-hand sides.

The new two-grid deflated methods are not as efficient as multigrid when it works well, but the new methods can potentially be used in situations where multigrid is not effective. Additionally, they can be combined with multigrid preconditioning.

Many facets of these two-grid deflated methods could use further investigation. For instance, three dimensional problems may have greater potential because the coarse grid matrix can be relatively smaller compared to the fine grid matrix. Other possible future work is to use more grid levels; see [26] for a multiple grid method for computing eigenvalues.

REFERENCES

- A. M. Abdel-Rehim, R. B. Morgan, and W. Wilcox. Deflated BiCGStab for linear equations in QCD problems. *Proceedings of Science, LAT2007*, pages 026/1–026/7, 2007.
- [2] A. M. Abdel-Rehim, A. Stathopoulos, and K. Orginos. Extending the eigCG algorithm to non-symmetric Lanczos for linear systems with multiple right-hand sides. *Numer. Linear Algebra Appl.*, 21:473493, 2013.
- [3] J. Baglama, D. Calvetti, G. H. Golub, and L. Reichel. Adaptively preconditioned GMRES algorithms. SIAM J. Sci. Comput., 20:243-269, 1998.
- [4] A. H. Baker, E. R. Jessup, and Tz. V. Kolev. A simple strategy for varying the restart parameter in GMRES(m). J. Comput. Appl. Math., 230:751–761, 2009.
- [5] A. H. Baker, E. R. Jessup, and T. Manteuffel. A technique for accelerating the convergence of restarted GMRES. SIAM J. Matrix Anal. Appl., 26:962–984, 2005.
- [6] A. Brandt. Multi-level adaptive solutions to boundary-value problems. Math. Comp., 31:333– 390, 1977.

- [7] S. Brenner and L. Scott. The Mathematical Theory of Finite Element Methods. Springer-Verlag, New York, NY, 1994.
- [8] W. L. Briggs, V. E. Henson, and S. F. McCormick. A Multigrid Tutorial, 2nd Edition. SIAM, Philadelphia, PA, USA, 2000.
- K. Burrage and J. Erhel. On the performance of various adaptive preconditioned GMRES strategies. Numer. Linear Algebra Appl., 5:101–121, 1998.
- [10] A. Chapman and Y. Saad. Deflated and augmented Krylov subspace techniques. Numer. Linear Algebra Appl., 4:43–66, 1997.
- [11] H. C. Elman, O. G. Ernst, and D. P. O'Leary. A multigrid method enhanced by Krylov subspace iteration for discrete Helmholtz equations. SIAM J. Sci. Comput., 23:1291–1315, 2001.
- [12] M. Embree. The tortoise and the hare restart GMRES. SIAM Review, 45:259–266, 2003.
- [13] Y. A. Erlangga and R. Nabben. Multilevel projection-based nested Krylov iteration for boundary value problems. SIAM J. Sci. Comput., 30:1572–1595, 2008.
- [14] Y. A. Erlangga, C. W. Oosterlee, and C. Vuik. A novel multigrid based preconditioner for heterogeneous Helmholtz problems. SIAM J. Sci. Comput., 27:1471–1492, 2006.
- [15] R. P. Fedorenko. The speed of convergence of one iterative process. USSR Comput. Math. Math. Phys., 4:227–235, 1964.
- [16] J. Frank and C. Vuik. On the construction of deflation-based preconditioners. SIAM J. Sci. Statist. Comput., 23:442–462, 2001.
- [17] A. Gaul, M. H. Gutknecht, J. Liesen, and R. Nabben. A framework for deflated and augmented Krylov subspace methods. SIAM J. Matrix Anal. Appl., 34:495–518, 2013.
- [18] L. Giraud, S. Gratton, X. Pinel, and X. Vasseur. Flexible GMRES with deflated restarting. SIAM J. Sci. Comput., 32:1858–1878, 2010.
- [19] S. A. Kharchenko and A. Y. Yeremin. Eigenvalue translation based preconditioners for the GMRES(k) method. Numer. Linear Algebra Appl., 2:51–77, 1995.
- [20] A. L. Laird and M. B. Giles. Preconditioned iterative solution of the 2D Helmholtz equation. Technical Report 02/12, Oxford Computer Laboratory, Oxford, UK, 2002.
- [21] R. B. Morgan. A restarted GMRES method augmented with eigenvectors. SIAM J. Matrix Anal. Appl., 16:1154–1171, 1995.
- [22] R. B. Morgan. On restarting the Arnoldi method for large nonsymmetric eigenvalue problems. Math. Comp., 65:1213–1230, 1996.
- [23] R. B. Morgan. GMRES with deflated restarting. SIAM J. Sci. Comput., 24:20–37, 2002.
- [24] R. B. Morgan and D. A. Nicely. Restarting the nonsymmetric Lanczos algorithm for eigenvalues and linear equations including multiple right-hand sides. SIAM J. Sci. Comput., 33:3037– 3056, 2011.
- [25] R. B. Morgan and W. Wilcox. Deflated iterative methods for linear equations with multiple right-hand sides. arXiv:math-ph/0405053v2, 2004.
- [26] R. B. Morgan and Z. Yang. Multigrid Arnoldi for eigenvalues. technical report, 2016.
- [27] M. L. Parks, E. de Sturler, G. Mackey, D. D. Johnson, and S. Maiti. Recycling Krylov subspaces for sequences of linear systems. SIAM J. Sci. Comput., 28:1651–1674, 2006.
- [28] J. Sifuentes. Preconditioned iterative methods for inhomogeneous acoustic scattering applications. PhD Thesis, Rice University, 2010.
- [29] J. A. Sifuentes, M. Embree, and R. B. Morgan. GMRES convergence for perturbed coefficient matrices, with application to approximate deflation preconditioning. SIAM J. Matrix Anal. Appl., 34:1066–1088, 2013.
- [30] P. Sonneveld and M. B. van Gijzen. IDR(s): a family of simple and fast algorithms for solving large nonsymmetric linear systems. SIAM J. Sci. Comput., 31:1035–1062, 2008.
- [31] A. Stathopoulos and K. Orginos. Computing and deflating eigenvalues while solving multiple right hand side linear systems in quantum chromodynamics. SIAM J. Sci. Comput., 32:439–462, 2010.
- [32] U. Trottenberg, C. W. Oosterlee, and A. Schuller. Multigrid. Academic Press, London, 2001.
- [33] H. A. van der Vorst. Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. SIAM J. Sci. Stat. Comput., 13:631–644, 1992.
- [34] M. B. van Gijzen and P. Sonneveld. Algorithm 913: An elegant IDR(s) variant that efficiently exploits bi-orthogonality properties. ACM Trans. Math. Soft., 38:5:1–5:19, 2011.
- [35] B. Zhong and R. B. Morgan. Complementary cycles of restarted GMRES. Numer. Linear Algebra Appl., 15:559–571, 2008.