

DEFLATED GMRES FOR SYSTEMS WITH MULTIPLE SHIFTS AND MULTIPLE RIGHT-HAND SIDES*

DEAN DARNELL[†], RONALD B. MORGAN[‡], AND WALTER WILCOX[§]

Abstract. We consider solution of multiply shifted systems of nonsymmetric linear equations, possibly also with multiple right-hand sides. First, for a single right-hand side, the matrix is shifted by several multiples of the identity. Such problems arise in a number of applications, including lattice quantum chromodynamics where the matrices are complex and non-Hermitian. Some Krylov iterative methods such as GMRES and BiCGStab have been used to solve multiply shifted systems for about the cost of solving just one system. Restarted GMRES can be improved by deflating eigenvalues for matrices that have a few small eigenvalues. We show that a particular deflated method, GMRES-DR, can be applied to multiply shifted systems.

In quantum chromodynamics, it is common to have multiple right-hand sides with multiple shifts for each right-hand side. We develop a method that efficiently solves the multiple right-hand sides by using a deflated version of GMRES and yet keeps costs for all of the multiply shifted systems close to those for one shift. An example is given showing this can be extremely effective with a quantum chromodynamics matrix.

Key words. linear equations, GMRES, BiCG, deflation, eigenvalues, QCD, multiple shifts, multiple right-hand sides

AMS subject classifications. 65F10, 15A06

1. Introduction. We consider the iterative solution of a large system of linear equations that not only has multiple right-hand sides, but also has multiple shifts for each right-hand side. Let $nrhs$ be the number of right-hand sides and ns be the number of shifts. Then the problem is

$$(1.1) \quad (A - \sigma_i I)x_i^j = b^j,$$

with $j = 1, \dots, nrhs$ and $i = 1, \dots, ns$. Here A is a large matrix which may be nonsymmetric or complex non-Hermitian. We assume there is no preconditioning. The shift σ_1 will be referred to as the base shift. For the rest of the paper, we will leave off the superscripts when it is clear which right-hand side is being dealt with.

Systems with multiple right-hand sides occur in many applications (see, for example, [18] for some applications). There are several applications that need solution of multiply shifted systems, for example, control theory [8], and time-dependent differential equations and frequency response computations [17]. Some important problems in lattice quantum chromodynamics (lattice QCD) have both multiple right-hand sides and multiple shifts. For example, the Wilson-Dirac formulation [7] and overlap fermion [34, 1] computations both lead to such problems. Very large complex non-Hermitian matrices are needed. For Wilson-Dirac matrices, the right-hand sides represent different noise vectors. The shifts correspond to different quark masses that are used in an extrapolation.

*This work was partially supported by the National Science Foundation, Computational Mathematics Program under grant 0310573. The second author was also supported by the Baylor University Sabbatical Program.

[†]Department of Physics, Baylor University, Waco, TX 76798-7316. (Dean_Darnell@baylor.edu).

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

[§]Department of Physics, Baylor University, Waco, TX 76798-7316. (Walter_Wilcox@baylor.edu).

A standard way to solve systems with multiple right-hand sides is to use a block approach [38, 43, 29]. However, block methods can be costly due to the orthogonalization expense, and also storage requirements can be prohibitive. Block methods are not generally used in lattice QCD. Even if multiple right-hand sides are solved sequentially instead of together in a block, it is important to take advantage of the fact that several systems share the same matrix. Information developed during the solution of one right-hand side can be used to help with others. This idea has been developed in different ways, including in seed methods [5] and with Richardson iteration [42]. Deflation methods can also be used (see [28, 37] and their references). Deflation approaches appearing in QCD literature include [9, 12, 11, 35]. Deflation involves computing eigenvectors corresponding to the smallest eigenvalues (or other outstanding eigenvalues) and using them to remove the eigenvalues from the effective spectrum for the iterative method. GMRES-DR [28] is a deflation method that has been adapted for multiple right-hand sides [31]. Some details of GMRES-DR are given in the next section.

Krylov methods have been developed for shifted matrix problems. Again, some details are in the next section.

The goal of this paper is to develop deflated methods for dealing with multiple shifts along with both a single right-hand side and multiple right-hand sides. Section 2 has background material on current methods for solving shifted systems and on deflated methods for multiple right-hand sides. In Section 3, deflation is used for multiply shifted systems with a single right-hand side, or the first of several right-hand sides. A multiply-shifted version of GMRES-DR is developed. There is also some comparison of GMRES with FOM for multiple shifts. Section 4 looks at solving the second and subsequent right-hand sides using eigenvector information and GMRES. Then the next section deals with the case of related right-hand sides.

2. Review.

2.1. Krylov methods for shifted systems. The Krylov subspace generated with the shifted matrix $A - \sigma_i I$ and with starting vector r_0 is

$$(2.1) \quad \text{Span}\{r_0, (A - \sigma_i I)r_0, (A - \sigma_i I)^2 r_0, \dots, (A - \sigma_i I)^{m-1} r_0\}.$$

Krylov subspaces are shift invariant in that the subspace is the same regardless of the choice of σ_i . Therefore, one Krylov subspace can be used to solve several shifted systems as long as all the systems have the same right-hand side, or at least right-hand sides that are multiples of each other. Note that we are assuming there is no preconditioning. With a preconditioner, systems with different shifts would no longer have equivalent Krylov subspaces. So for the non-preconditioned case, it is fairly straightforward to develop versions of nonrestarted Krylov methods for multiply shifted systems. Such versions have been given for the conjugate gradient method [15, 47], nonrestarted GMRES [8], and both QMR and TFQMR [17]. A multiply shifted version of BiCGStab has also been developed [19].

Restarted Krylov methods are not as straightforward. After a restart, all systems need to have parallel right-hand sides. This means that all residual vectors formed at the end of a cycle of the Krylov method need to be multiples of each other. For the FOM method [38], this happens automatically [41]. For GMRES [40, 38], the residuals can be forced to be parallel, as shown by Frommer and Glässner [20] (see also [14]). To see how this is possible, we need the Arnoldi recurrence [38]:

$$(2.2) \quad AV_m = V_{m+1}\bar{H}_m,$$

where V_m is a n by m matrix whose columns span the Krylov subspace, V_{m+1} is the same except for an extra column and \bar{H}_m is an upper-Hessenberg $m+1$ by m matrix. For a shifted system, this becomes

$$(2.3) \quad (A - \sigma_i I)V_m = V_{m+1}(\bar{H}_m - \sigma_i \bar{I}),$$

where \bar{I} is the $m+1$ by m portion of an identity matrix.

Say that after a restart the approximate solution for the i th shifted system is $\tilde{x}_{0,i}$ and the residual vector is $r_{0,i} = \beta_i r_{0,1}$. Here the first subscript of 0 on x and r shows they are the vectors from the beginning of the cycle, and the second subscript tells which shift. Note it is assumed that the residual is parallel to that of the base shift. The system is then

$$(A - \sigma_i I)(x_i - \tilde{x}_{0,i}) = r_{0,i} = \beta_i r_{0,1}.$$

For the next cycle, let the base system have standard GMRES solution $V_m d_1$, where d_1 is the solution of $\min \|c - (\bar{H}_m - \sigma_1 \bar{I})d_1\|$, with $r_{0,1} = V_{m+1}c$. Let r_i be the new residual vector for the i th shifted system after this cycle. We need the approximate solution $V_m d_i$ to be chosen so that $r_i = \beta_i^{new} r_1$, for some scalar β_i^{new} . So

$$(2.4) \quad \beta_i r_{0,1} - (A - \sigma_i I)V_m d_i = \beta_i^{new} (r_{0,1} - (A - \sigma_1 I)V_m d_1).$$

After using the shifted Arnoldi recurrence (2.3), we have

$$(2.5) \quad V_{m+1}(\beta_i c - (\bar{H}_m - \sigma_i \bar{I})d_i) = \beta_i^{new} V_{m+1}(c - (\bar{H}_m - \sigma_1 \bar{I})d_1).$$

Next the V_{m+1} can be dropped, and we let $s = c - (\bar{H}_m - \sigma_1 \bar{I})d_1$ and rearrange:

$$(\bar{H}_m - \sigma_i \bar{I})d_i = \beta_i c - \beta_i^{new} s.$$

We now use a QR factorization, $\bar{H}_m - \sigma_i \bar{I} = QR$, with Q being an $m+1$ by $m+1$ orthogonal matrix (or unitary in the complex case) and R being $m+1$ by m upper-triangular, and get

$$Rd_i = \beta_i Q^H c - \beta_i^{new} Q^H s.$$

The value of β_i^{new} can be determined from the last row (note the left side of the equation has zero in the last row). Then solution of an upper-triangular system determines d_i .

A shortcut formula for the residual norm of the i th shifted system is

$$(2.6) \quad \|r_i\| = \|\beta_i c - (\bar{H}_m - \sigma_i \bar{I})d_i\|.$$

The new approximate solution is $\tilde{x}_i = \tilde{x}_{0,i} + V_m d_i$. The systems then become $A(x_i - \tilde{x}_i) = \beta_i^{new} r_1$. For more details, see [20]. We will refer to this approach as GMRES-Shifts or GMRES-Sh. Note that only the base system has the minimum residual property. The solution of the other shifted systems is not equivalent to GMRES applied to those systems. This is in contrast to FOM-Shifts [41], which has each shifted system solved with the FOM approach. Associated with Krylov methods such as GMRES and FOM is a polynomial. In fact, the new approximate solution at the end of a cycle can be written as $\tilde{x}_i = \tilde{x}_{0,i} + p(A)r_0$, where p is a polynomial of degree $m-1$ or less. Like FOM-Sh, GMRES-Sh uses the same polynomial (after shifting) for all shifted systems [20, 41]. If true GMRES was applied to each of the systems, they would have different polynomials.

2.2. Deflated GMRES. Small subspaces for restarted GMRES can slow convergence for difficult problems. Deflated versions of restarted GMRES [25, 23, 13, 6, 39, 2, 3, 4, 10, 27, 28] can improve this, when the problem is difficult due to a few small eigenvalues. One of these approaches is related to Sorensen's implicitly restarted Arnoldi method for eigenvalues [44] and is called GMRES with implicit restarting [27]. A mathematically equivalent method, called GMRES with deflated restarting (GMRES-DR) [28], is also related to Wu and Simon's restarted symmetric Lanczos eigenvalue method [46]. See [26, 45, 33] for some other related eigenvalue methods.

We will concentrate on GMRES-DR [28], because it is efficient and relatively simple. Approximate eigenvectors corresponding to the small eigenvalues are computed at the end of each cycle and are put at the beginning of the next subspace. Letting r_0 be the initial residual for the linear equations at the start of the new cycle and $\tilde{y}_1, \dots, \tilde{y}_k$ be harmonic Ritz vectors [24, 16, 36, 32], the subspace of dimension m used for the new cycle of GMRES-DR(m, k) is

$$(2.7) \quad \text{Span}\{\tilde{y}_1, \tilde{y}_2, \dots, \tilde{y}_k, r_0, Ar_0, A^2r_0, A^3r_0, \dots, A^{m-k-1}r_0\}.$$

This can be viewed as a Krylov subspace generated with starting vector r_0 augmented with approximate eigenvectors. Remarkably, the whole subspace turns out to be a Krylov subspace itself (though not with r_0 as starting vector) [27]. Once the approximate eigenvectors are moderately accurate, their inclusion in the subspace for GMRES essentially deflates the corresponding eigenvalues from the linear equations problem.

GMRES-DR generates a recurrence similar to the Arnoldi recurrence (2.2). It is

$$(2.8) \quad AV_m = V_{m+1}\bar{H}_m,$$

where V_m is a n by m matrix whose columns span the subspace (2.7), V_{m+1} is the same except for an extra column and \bar{H}_m is an $m+1$ by m matrix that is upper-Hessenberg except for a full $k+1$ by k leading portion. A part of this recurrence can be separated out to give

$$(2.9) \quad AV_k = V_{k+1}\bar{H}_k,$$

where V_k is a n by k matrix whose columns span the subspace of approximate eigenvectors, V_{k+1} is the same except for an extra column and \bar{H}_k is a full $k+1$ by k matrix. This recurrence allows access to both the approximate eigenvectors and their products with A while requiring storage of only $k+1$ vectors of length n . The approximate eigenvectors in GMRES-DR actually span a small Krylov subspace of dimension k .

2.3. Deflated GMRES for multiple right-hand sides. The multiple right-hand side approach that will be adapted here for multiple shifts is called GMRES-Proj. The first right-hand side is solved with GMRES-DR, then the eigenvector information thus generated is used to deflate eigenvalues from other right-hand sides [31]. This is done with a simple minimum residual projection over the eigenvectors alternated with cycles of regular GMRES(m). The expense of the projection does not generally add much to the GMRES cost.

The algorithm for the projection is given next. It projects over the space of harmonic Ritz vectors spanned by the columns of V_k in Equation (2.9). This requires only $3k+2$ vector operations of length n .

Minres Projection for V_k

1. Let the current approximate solution be \tilde{x}_0 and the current system of equations be $A(x - \tilde{x}_0) = r_0$. Let V_{k+1} and \bar{H}_k be from Equation (2.9).
2. Solve $\min \|c - \bar{H}_k d\|$, where $c = (V_{k+1})^H r_0$.
3. The new approximate solution is $\tilde{x} = \tilde{x}_0 + V_k d$.
4. The new residual vector is $r = r_0 - AV_k d = r_0 - V_{k+1} \bar{H}_k d$.

The GMRES-Proj method that follows is for all right-hand sides except for the first one. See [31] for more details.

GMRES(m)-Proj(k)

1. After applying the initial guess \tilde{x}_0 , let the system of equations be $A(x - \tilde{x}_0) = r_0$.
2. If it is known that the right-hand sides are closely related, project over the previous computed solution vectors.
3. Apply the Minres Projection for V_k . This uses the V_{k+1} and \bar{H}_k matrices developed while solving the first right-hand side with GMRES-DR.
4. Apply one cycle of GMRES(m).
5. Test the residual norm for convergence (can also test during the GMRES cycles). If not satisfied, go back to Step 3.

3. Deflated GMRES with Multiple Shifts. Subsection 2.1 gave some details of how restarted GMRES can be implemented in order to simultaneously solve multiply shifted systems. For the deflated restarted GMRES method GMRES-DR, we can also solve multiply shifted systems concurrently. The key is that GMRES-DR has subspaces that are Krylov subspaces (as mentioned in Section 2.2), and they always contain the current right-hand side vector. The derivation in Subsection 2.1 for solving multiply shifted systems with GMRES also applies here with very slight change. Because GMRES-DR has Krylov subspaces, it has the Arnoldi-like recurrence (2.8) which can be used in place of GMRES's Arnoldi recurrence (2.2) (\bar{H}_m is not upper-Hessenberg, but this does not affect the derivation). Also, since the subspaces contain the right-hand side, $r_{0,i}$ is again $V_{m+1}(\beta_i c)$ for some c , but here c is not a multiple of e_1 . Next is the algorithm for the shifting part of the new method that we call GMRES-DR with shifts or GMRES-DR-Sh. It is the same as for GMRES-Sh, except it uses GMRES-DR instead of GMRES and has the different form of \bar{H}_m . See [28] for more on the GMRES-DR portion. See [21] for a multi-shifted version of the related method GMRES-E [25].

GMRES-DR(m,k)-Sh

1. At the beginning of a cycle of GMRES-DR-Sh, assume the current problem is $(A - \sigma_i I)(x_i - \tilde{x}_{0,i}) = \beta_i r_{0,i}$, with $\beta_1 = 1$, and where $\tilde{x}_{0,i}$ is the current approximate solution to the i th shifted system.
2. Apply GMRES-DR(m,k) to A and generate Equation (2.8): $AV_m = V_{m+1} \bar{H}_m$.
3. For the base system, solve the minimum residual reduced problem $\min \|c - (\bar{H}_m - \sigma_1 \bar{I}) d_1\|$ for d_1 , where $c = V_{m+1}^H r_{0,1}$ and \bar{I} is the first m columns of the $m + 1$ by $m + 1$ identity matrix. The new approximate solution is $\tilde{x}_1 = \tilde{x}_{0,1} + V_m d_1$. The new residual vector is $r_1 = r_{0,1} - AV_m d_1 = r_{0,1} - V_{m+1} \bar{H}_m d_1$.
4. For the other shifted systems $i = 2, \dots, ns$, form $s = c - (\bar{H}_m - \sigma_1 \bar{I}) d_1$. Apply a QR factorization: $\bar{H}_m - \sigma_i \bar{I} = QR$. Solve $Rd_i = \beta_i Q^H c - \beta_i^{new} Q^H s$, using the last row to solve for β_i^{new} and the first m rows for d_i .

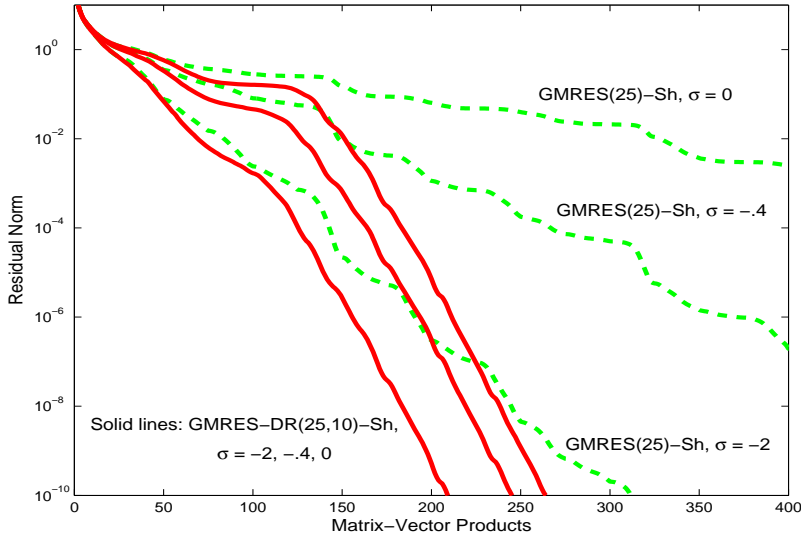


FIG. 3.1. Solution of a system with multiple shifts.

5. The new approximate solution of the i th system is $\tilde{x}_i = x_{0,i} + V_m d_i$, and the new residual is $r_i = \beta_i^{new} r_1$.
6. Test the residual norms for convergence. If not satisfied, for $i = 2 \dots ns$, set $\beta_i = \beta_i^{new}$ and for $i = 1 \dots ns$, set $\tilde{x}_{0,i} = \tilde{x}_i$ and $r_{0,i} = r_i$. Then go back to Step 1.

The expense for GMRES-DR(m,k)-Sh is only a little greater than for a single shift. Counting only length n vector operations (such as dot products and daxpys), one cycle requires $m - k$ matrix-vector products and about $m^2 + (k + 3)m - k^2 + (ns - 1)m$ vector operations. The last term is for forming the new approximate solutions for the $ns - 1$ shifts beyond a single shift.

Example 1. We test the shifted version of deflated GMRES and compare it to shifted regular GMRES. All experiments are done in Matlab 6 on a Dell workstation running Windows XP, except for example 6. The Matlab command “eig” is used to compute eigenvectors, again except for example 6. The matrix has $n = 1000$ and is bidiagonal with $0.1, 1, 2, 3, \dots, 998, 999$ on the main diagonal and 1’s on the super-diagonal. The right-hand side is generated randomly. The shifts are $\sigma = 0, -0.4, -2$. GMRES(25)-Sh is compared with GMRES-DR(25,10)-Sh. The results are given in Figure 3.1. This problem has small eigenvalues which slow down restarted GMRES, particularly for the base system with $\sigma = 0$. Shifting the matrix even by just 0.4 improves convergence of GMRES(25), mainly because the smallest eigenvalue is moved from 0.1 to 0.5. GMRES-DR converges very rapidly once it generates approximations to the eigenvectors corresponding to these eigenvalues. The convergence of GMRES-DR for all three shifted systems is similar, since once the small eigenvalues are essentially removed by the deflation, shifting does not have such an important effect.

In Example 1, the second and third shifted systems converge faster than the first system. However, in some situations, there can be convergence problems for the non-

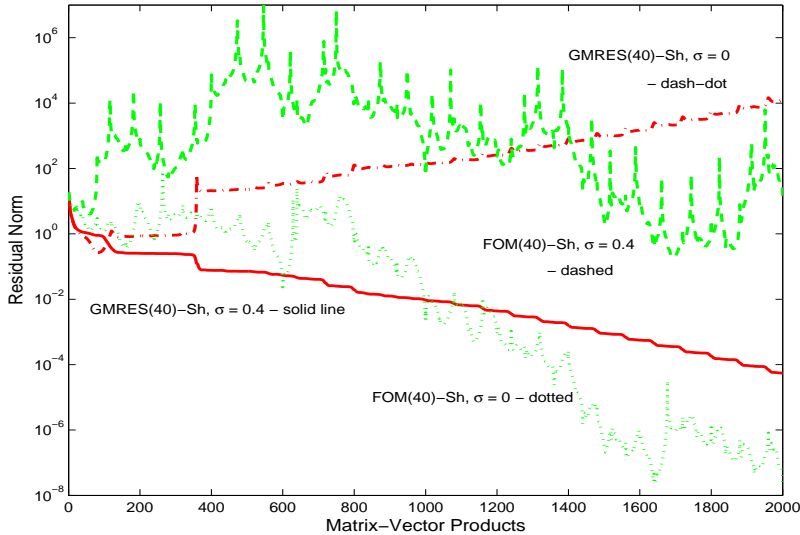
base systems [41]. Simoncini [41] compares multiply shifted GMRES and FOM. Since FOM automatically has residuals parallel for all shifted systems, it can be argued that it is a more natural approach [41]. However, it is also a matter of where the roots of the FOM and GMRES polynomials fall in relation to the shift. The next example demonstrates this. As mentioned earlier, GMRES-Sh uses the same polynomial for all shifted systems [20]. Likewise, FOM-Sh sticks with one polynomial for all shifts. The regular GMRES polynomial for an unshifted matrix is scaled to be 1 at zero and needs to be small over the spectrum. For shifted systems, we have a choice of viewing the polynomial as being 1 at zero and the spectrum shifted or the polynomial being 1 at the shift and the spectrum fixed as that of A . We chose the latter. So for GMRES-Sh with the base system $A - \sigma_1 I$, we view the polynomial chosen by GMRES as being 1 at σ_1 and needing to be small over the spectrum of A . In the next two examples, plots are given of the roots of these polynomials. Because the polynomial needs to be somewhat small over the spectrum of A , a small value of the polynomial at the shift can cause a problem. With the normalizing, there may then be large values at eigenvalues of A . So for a Krylov method to be effective, the roots of the polynomials need to generally stay away from the shift.

Example 2. For the same matrix as in Example 1, we apply GMRES(40)-Sh and FOM(40)-Sh with shifts $\sigma = 0.4, 0$. The results are shown in Figure 3.2. For the base shift of $\sigma = 0.4$, GMRES works better, but FOM is more effective on the next shift. Figure 3.3 shows the harmonic Ritz values [24, 36] nearest the shift for 50 cycles of GMRES. These are the roots of the GMRES polynomial for the system $(A - 0.4I)x_1 = b$, shifted so they correspond with the spectrum of A . The harmonic Ritz values avoid the region around 0.4, which is good, since the shifted GMRES polynomial for $\sigma = 0.4$ needs to have value 1 at 0.4 and be somewhat small over the spectrum of A . This polynomial may not be effective if it has a root near to 0.4. GMRES(40) for $\sigma = 0.4$ is able to slowly converge for this fairly difficult, indefinite problem. Meanwhile, FOM(40)-Sh with $\sigma = 0.4$ does not converge, because as shown in Figure 3.3, the roots of the shifted FOM polynomials, which are the Ritz values, are not separated from 0.4. For the second shift value of 0, GMRES-Sh is not effective. Too many harmonic Ritz values occur not only near 0, but also on both sides of it. FOM gives erratic convergence due to some Ritz values near 0, but it does converge.

Although not included in Figure 3.2, we also tested a third shift $\sigma = -1$. GMRES-Sh again does not converge because of harmonic Ritz values near -1. This is in spite of the fact that GMRES would have no trouble if this was the base shift. The Ritz values for FOM do not occur near 1, so FOM converges to under 10^{-8} in 400 iterations. One can also construct an example where GMRES-Sh is better than FOM-Sh for a non-base shift by choosing the shift in a zone where GMRES-Sh has fewer harmonic Ritz values than FOM-Sh has Ritz values. However, the main point is that multi-shifted GMRES and FOM must both be used with some caution.

We continue the example by testing GMRES-DR(40,2)-Sh and FOM-DR(40,2)-Sh for the same shifts, $\sigma = 0.4, 0$. Figure 3.4 has the results. We observe three things: deflating just the two smallest eigenvalues gives much faster convergence, there is no longer much trouble from having harmonic Ritz values near the shifts, and GMRES-DR-Sh converges a little faster than FOM-DR-Sh.

Finally, it is suggested in [14] that alternation of which shift is chosen as the base shift can improve convergence. We test this idea on GMRES-DR(40,2)-Sh with shifts $\sigma = 0.4, 0$. After eight cycles, alternating shifts gives residual norms of $2.2e - 8$ and $3.0e - 10$ respectively for the two shifted systems. This compares to $2.2e - 7$

FIG. 3.2. *Shifted GMRES and FOM*

and $3.0e - 8$ in the test that was given in Figure 3.4. So alternating shifts does give some improvement for deflated GMRES. However, there is little difference when 10 eigenvalues are used. Deflating eigenvalues seems to make the choice of base shift not such an important issue. For the examples in the rest of this paper, the matrices and shifts are such that the base systems are the ones with the slowest convergence. In [20], Theorem 3.3 shows that for a positive real matrix with negative real shifts, the residuals of the non-base systems are smaller than the base ones. So for such problems, it is easy to choose which shift will give the slowest convergence. Example 6 has a QCD matrix with this property and the non-base shifted systems do converge more quickly.

4. Deflated GMRES for multiple right-hand sides and multiple shifts.

We now consider solving multiply shifted systems that also have multiple right-hand sides. It is important to reuse information or share information among the right-hand sides. It is possible to design multi-shifted versions of both Block-GMRES [38] and Block-GMRES-DR [29]. However, here we will concentrate on a non-block approach. The right-hand sides are solved separately, and eigenvector information from solution of the first right-hand side is used to assist the subsequent ones. More specifically, we will generalize for multiple shifts the GMRES-Proj approach mentioned in Section 2. See [31] for more on this method, including comparison with block methods.

First some of the difficulties of deflating for subsequent right-hand sides will be discussed. Suppose the first right-hand side has been solved and approximate eigenvectors have been generated. Then for the non-shifted case, there are several ways to deflate eigenvalues. Some of these are given in [31]. However, generally they do not work for multiply shifted systems. For example, some deflation approaches involve building a preconditioner from the approximate eigenvectors [2, 3, 31]. As mentioned earlier, shifted systems cannot be solved together if there is preconditioning.

For the GMRES-Proj method, there is trouble with one of the two phases. We

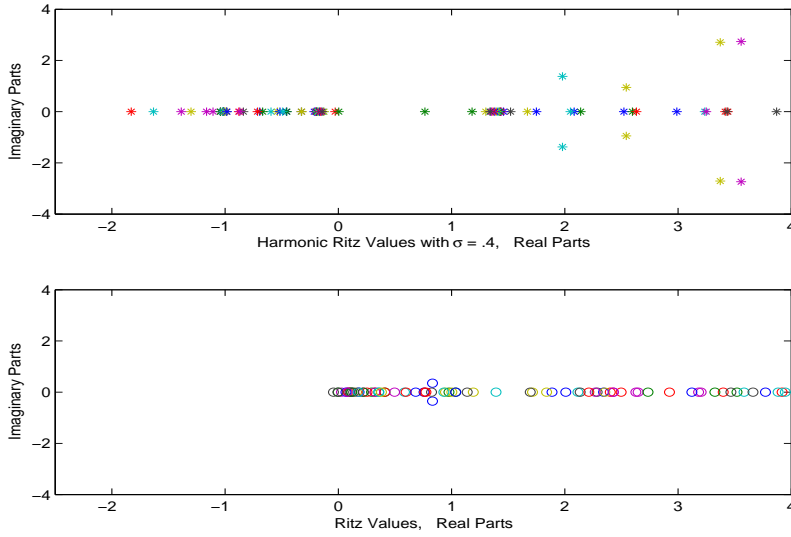


FIG. 3.3. *Distribution of Smallest Harmonic and Regular Ritz values, $m=40$, 50 cycles*

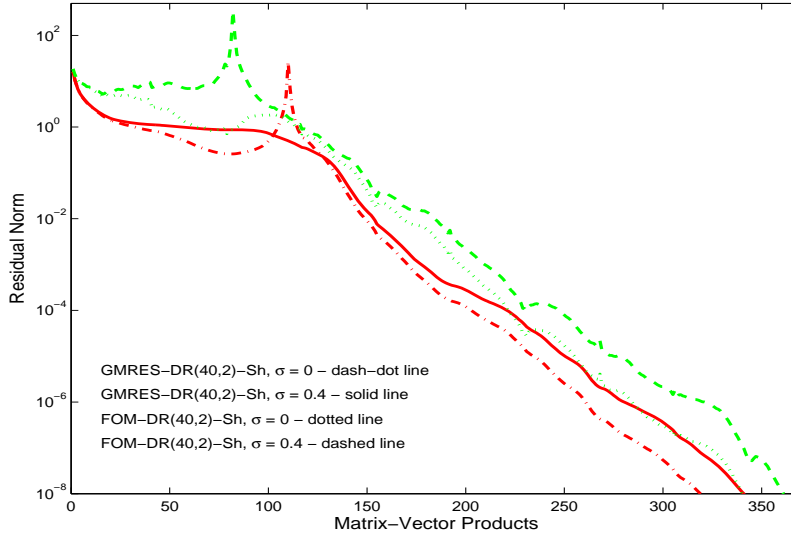


FIG. 3.4. *Shifted GMRES-DR and FOM-DR*

know the GMRES portion can be adapted to keep residuals parallel for multiple shifts. However, the phase with projection over approximate eigenvectors generally fails to produce parallel residual vectors. Even though this projection is over a Krylov subspace of dimension k spanned by the columns of V_k , this subspace does not contain the current right-hand side (the residual vector). So the derivation in Section 2.1 does not work with V_{k+1} and \bar{H}_k from (2.9) replacing V_{m+1} and \bar{H}_m from (2.2). Specifically, the transition from Equation (2.4) to Equation (2.5) is not possible since $r_{0,i}$ and $r_{0,1}$

are not in the span of the columns of V_{k+1} . One case where the projection does keep the residual vectors parallel is with exact eigenvectors. We now show this.

THEOREM 4.1. *Assume that before the minres projection, the shifted systems are*

$$(A - \sigma_i I)(x_i - \tilde{x}_{0,i}) = r_{0,i}$$

with $r_{0,i} = \beta_i r_{0,1}$ for $i = 2, \dots, ns$. Let z_1, z_2, \dots, z_k be eigenvectors of A . Then after the minres projection over the subspace $\text{Span}\{z_1, z_2, \dots, z_k\}$, the residual vectors are parallel.

Proof. Let Z_k be the matrix with z_1, \dots, z_k as columns. For exact eigenvectors, the minres projection is equivalent to Galerkin [38]. The Galerkin projection gives the reduced problem:

$$Z_k^H (A - \sigma_i I) Z_k d_i = \beta_i Z_k^H r_{0,1}.$$

Solving gives

$$d_i = \beta_i (\Lambda_k - \sigma_i I_k)^{-1} (Z_k^H Z_k)^{-1} Z_k^H r_{0,1},$$

where Λ_k is the k by k diagonal matrix with diagonal entries λ_1 through λ_k and I_k is the k by k identity matrix. The residual vector after projecting is then

$$\begin{aligned} r_i &= \beta_i r_{0,1} - (A - \sigma_i I) Z_k d_i \\ &= \beta_i r_{0,1} - \beta_i (A - \sigma_i I) Z_k (\Lambda_k - \sigma_i I_k)^{-1} (Z_k^H Z_k)^{-1} Z_k^H r_{0,1} \\ &= \beta_i r_{0,1} - \beta_i Z_k (\Lambda_k - \sigma_i I_k) (\Lambda_k - \sigma_i I_k)^{-1} (Z_k^H Z_k)^{-1} Z_k^H r_{0,1} \\ &= \beta_i (I - Z_k (Z_k^H Z_k)^{-1} Z_k^H) r_{0,1}. \end{aligned}$$

This shows that all r_i are multiples of each other. \square

So one option for GMRES-Proj with multiple shifts is to use only fairly accurate eigenvectors. We could sort through the approximate eigenvectors computed by GMRES-DR and apply only those with acceptable accuracy to the projection in GMRES-Proj. This is now tested.

Example 3. For the same matrix as in Example 1 and two randomly generated right-hand sides, we solve the $\sigma_1 = 0$ system with the first right-hand side to accuracy of relative residual norm below 10^{-10} . This takes 250 matrix-vector products. The second column of Table 4.1 shows the accuracy of the k th eigenvector thus produced. For example, the residual norm of the tenth approximate eigenvector is 4.1e-2, while the fourth is much better at 1.0e-6. Now we consider solution of the second right-hand side using the first k eigenvectors with k from 10 down to 2. The third column gives the number of matrix-vector products for the relative residual norm of the base shifted system with the second right-hand side to reach 1.e-10. We see that convergence is better using all ten eigenvectors. However, the fourth column gives the accuracy attained by the worst of the last two shifted systems (usually it is the third shift). It only reaches residual norm of 3.3e-4 if all 10 approximate eigenvectors are used in the projection. With only four eigenvectors, the residual norm reaches a better level of 5.4e-8, but the convergence is almost twice as slow.

The last three columns repeat this information for the case of solving the first right-hand side system to greater accuracy of relative residual norm of 1.e-14 (385 matrix-vector products). The second right-hand side systems can then be solved more accurately. Even with $k = 8$ eigenvectors, all shifted systems reach accuracy of 1.3e-9 or better compared to 3.7e-5 for the previous case of less accurate eigenvectors.

TABLE 4.1
Effect of projecting over different accuracies of eigenvectors

	250 mvp's for 1st		385 mvp's 1st			
k	eig. res.	mvp's	lin. eqs. res.	eig. res.	mvp's	lin. eqs. res.
10	4.1e-2	135	3.3e-4	4.4e-1	135	1.7e-6
8	3.3e-3	165	3.7e-5	7.1e-6	165	1.3e-9
6	8.0e-5	180	2.8e-6	3.3e-10	180	1.3e-9
4	1.0e-6	255	5.4e-8	2.4e-11	255	5.5e-10
2	5.4e-9	435	1.5e-8	9.9e-12	435	3.6e-10

The problem with this approach is that the eigenvector computation during solution of the first right-hand side needs to be done to considerable accuracy, since we do not want to slow down convergence of the subsequent systems. If many right-hand sides are to be solved, this extra expense might not be significant. However, we next propose an approach without this concern of needing accurate eigenvectors.

The key idea is that although the residual vectors cannot be kept parallel, they can be chosen so that they relate to each other. We force the residuals of the non-base systems to be parallel to the residual of the base system except for a component in the direction of v_{k+1} , the last column of the V_{k+1} matrix from Equation (2.9). We then continue solving but ignore this component. At the end, a correction can be done. However, for this correction, we need solution of shifted systems with one additional right-hand side, namely v_{k+1} . We begin discussion of this approach with the aspect of keeping residuals parallel except for the one component, then move on to the correction phase.

Suppose we have shifted systems with parallel residuals, namely $(A - \sigma_i I)(x_i - \tilde{x}_{0,i}) = \beta_i r_{0,1}$, with $\beta_1 = 1$. A projection over the columns of V_k can be implemented to give the desired relationship between residuals. First apply minres projection over the approximate eigenvectors spanned by the columns of V_k to the base system. The base system residual vector is then $r_1 = r_{0,1} - (A - \sigma_1 I)V_k d_1$. We need d_i to be chosen so that

$$(4.1) \quad r_i = \beta_i r_1 + \gamma_i v_{k+1},$$

for some scalar γ_i . So we need

$$r_{0,i} - (A - \sigma_i I)V_k d_i = \beta_i(r_{0,1} - (A - \sigma_1 I)V_k d_1) + \gamma_i v_{k+1}.$$

After using that $r_{0,i} = \beta_i r_{0,1}$ and using the key recurrence for the approximate eigenvectors (2.9), we have

$$V_{k+1}(\bar{H}_k - \sigma_i \bar{I})d_i = \beta_i V_{k+1}(\bar{H}_k - \sigma_1 \bar{I})d_1 - \gamma_i v_{k+1}.$$

Next the V_{k+1} can be dropped, and we let $t = \beta_i(\bar{H}_k - \sigma_1 \bar{I})d_1$:

$$(4.2) \quad (\bar{H}_k - \sigma_i \bar{I})d_i = t - \gamma_i e_{k+1},$$

where e_{k+1} is the $k + 1$ st coordinate vector of length $k + 1$. We ignore the last row of (4.2) and solve the first k equations for the unknown vector d_i of length k . Then Equation (4.2) is automatically true for some γ_i . So we have (4.1). Note the β_i 's do not change during this projection over approximate eigenvectors, unlike in the GMRES-Sh portion.

Now we will look at the correction phase that is needed at the end of GMRES-*Proj*. Assume that we have already solved shifted systems with the extra right-hand side v_{k+1} (this solution will be discussed next) and have

$$(4.3) \quad (A - \sigma_i I) s_i = v_{k+1}.$$

We assume that for a particular right hand side, the systems have been solved by GMRES-*Proj* to the point that the residual is only in the direction of v_{k+1} and the system is recast as

$$(4.4) \quad (A - \sigma_i I)(x_i - \tilde{x}_i) = r_i = \gamma_i v_{k+1},$$

for $i = 2, \dots, ns$ and for some scalar γ_i . Here \tilde{x}_i is the approximate solution to x_i . We perform a Galerkin projection for system (4.4) over the subspace spanned by the single vector s_i from solution of (4.3):

$$s_i^H (A - \sigma_i I) s_i \delta = \gamma_i s_i^H v_{k+1}.$$

Using Equation (4.3), this becomes

$$s_i^H v_{k+1} \delta = \gamma_i s_i^H v_{k+1}.$$

Then $\delta = \gamma_i$. To determine γ_i , we start with $r_i = \gamma_i v_{k+1}$. Multiplying both sides by v_{k+1}^H and using that v_{k+1} is of unit length gives

$$(4.5) \quad \gamma_i = v_{k+1}^H r_i.$$

The corrected solution of the system is $x_i = \tilde{x}_i + \delta s_i = \tilde{x}_i + v_{k+1}^H r_i s_i$.

We need to fill in the method for solving (4.3), the shifted systems with the extra right-hand side v_{k+1} . First, GMRES-*Proj* is applied. This includes projection over the approximate eigenvectors as described above alternating with cycles of GMRES. This continues until the residual is negligible except in the direction of v_{k+1} . So for the correction, we assume that

$$(4.6) \quad (A - \sigma_i I) \tilde{s}_i = r$$

is the current system, where

$$r = v_{k+1} - (A - \sigma_i I) \tilde{s}_i = \gamma_i v_{k+1},$$

for some scalar γ_i . Rearranging this residual equation gives

$$(4.7) \quad (A - \sigma_i I) \tilde{s}_i = (1 - \gamma_i) v_{k+1}.$$

Applying Galerkin projection over the subspace spanned by the single vector \tilde{s}_i to the system (4.6) gives

$$\tilde{s}_i^H (A - \sigma_i I) \tilde{s}_i \delta = \gamma_i \tilde{s}_i^H v_{k+1}.$$

With Equation (4.7), this becomes

$$(1 - \gamma_i) \tilde{s}_i^H v_{k+1} \delta = \gamma_i \tilde{s}_i^H v_{k+1},$$

and this simplifies to

$$\delta = \frac{\gamma_i}{1 - \gamma_i}.$$

So the corrected solution is $s_i = \tilde{s}_i + \frac{\gamma_i}{1 - \gamma_i} \tilde{s}_i = \frac{1}{1 - \gamma_i} \tilde{s}_i$. Finally, the γ_i can be determined to be $\gamma_i = v_{k+1}^H r_i$ as it was for (4.5).

We next list the algorithms. First is the solution of the systems with the extra right-hand side. Then the algorithm is given for the solution of the systems with the second and subsequent right-hand sides. They differ in the final correction phase and in that there may be solutions of related right-hand sides that can be used to help solution of the second and subsequent right-hand sides.

GMRES(m)-Proj(k)-Sh for the extra right-hand side v_{k+1}

1. Consider the systems with right-hand side v_{k+1} (and with all ns shifts).
2. At the beginning of a cycle of GMRES-Proj-Sh, assume the current problem is $(A - \sigma_i I)(x_i - \tilde{x}_{0,i}) = \beta_i r_{0,i}$, with $\beta_1 = 1$, and where $\tilde{x}_{0,i}$ is the current approximate solution to the i th shifted system.
3. Apply the Minres Projection for V_k to the base shift. This uses the V_{k+1} and \tilde{H}_k matrices developed while solving the first right-hand side with GMRES-DR-Sh.
4. For shifted systems $i = 2 \dots ns$, solve $(H_k - \sigma_i I)d_i = \beta_i(H_k - \sigma_1 I)d_1$, where H_k is the k by k portion of \tilde{H}_k . Set the new approximate solution as $\tilde{x}_i = \tilde{x}_{0,i} + V_k d_i$.
5. Apply one cycle of GMRES(m)-Sh.
6. Test the residual norms for convergence (can also test during the GMRES cycles). For the non-base systems, ignore the error term in the direction of v_{k+1} . Residual formula (2.6) can be used. If not satisfied, go back to Step 3. Otherwise conclude with Step 7.
7. Correction phase: Suppose the computed solution for the i th shifted system so far is \tilde{s}_i . The corrected solution is $s_i = (\frac{1}{1 - \gamma_i})\tilde{s}_i$, with $\gamma_i = v_{k+1}^H r$.

GMRES(m)-Proj(k)-Sh for the second and subsequent right-hand sides

The algorithm is the same as for previous algorithm except for

1. Consider the systems with the j th right-hand side (and with all ns shifts). If it is known that the right-hand sides are closely related, project over the previous computed solution vectors (see section 5).
7. Correction phase: Suppose the computed solution so far for the i th shifted system is \tilde{x}_i . Let the solution to the system with the extra right hand side v_{k+1} and shift σ_i be s_i . The corrected solution is $x_i = \tilde{x}_i + (v_{k+1}^H r)s_i$ and the corrected residual norm can be calculated.

We give the expense of a cycle of GMRES(m)-Proj(k)-Sh for either the extra right-hand side or any of the right-hand sides after the first. Required are m matrix-vector products and about $m^2 + 2m + (ns - 1)m$ length n vector operations for the GMRES-Sh part plus $3k + 2 + (ns - 1)k$ vector operations for the projection. There is no need for m here to be the same as the m in GMRES-DR-Sh. In the examples, m is set to the value of GMRES-DR-Sh's $m - k$.

Example 4. We use the same test matrix. All right-hand sides are generated randomly. The systems with the first right-hand side are solved with GMRES-DR(25,10)-Sh as in Example 1. Then the extra right-hand v_{k+1} systems are solved (for all

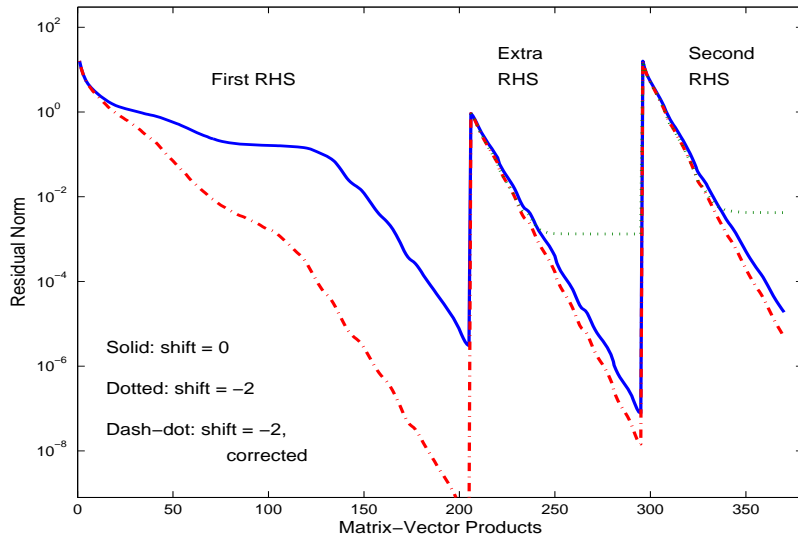


FIG. 4.1. *Solution of first rhs, extra rhs and second rhs with two shifts.*

shifts) with GMRES(15)-Proj(10)-Sh. Finally, the second right-hand side systems are also solved with GMRES(15)-Proj(10)-Sh. All relative residual tolerances are $rtol = 1.e - 6$. Figure 4.1 has residual curves for only two shifts, the base shift of zero and $\sigma_2 = -2$. The solid line is the base shift. The dotted line shows the uncorrected residual norm for the second shift, while the dash-dot line has the second shift residuals if they are corrected (actually the correction needs to be done only once at the end). The uncorrected residual norm for the second right-hand side and second shift levels off at $4.e-3$, but this is fixed by the correction phase. The convergence is faster than for GMRES-DR-Sh, because the eigenvectors are used from the beginning to speed up the convergence. Also the cost of GMRES(15)-Proj(10)-Sh is less than for GMRES-DR(25,10)-Sh, because it is fairly inexpensive to project over the approximate eigenvectors compared to keeping the eigenvectors in the GMRES-DR subspace. Here the expense for the extra right-hand side is fairly significant, however it will not be if there are more right-hand sides. Figure 4.2 has the case of solving a total of 10 right-hand sides. Also, the extra right-hand side is solved only to relative residual tolerance of $1.e-3$. Now the expense for the extra right-hand side v_{k+1} is small compared to amount saved by speeding up solution of all the right-hand sides. We also note that the expense for solving all right-hand sides after the first is fairly constant. The method gets a little cheaper per right-hand side as more are solved, because the extra cost for the first right-hand side and for the v_{k+1} right-hand side becomes less significant.

Example 5. At the end of the previous example, the extra right-hand side is solved to low accuracy, but the correction for the subsequent right-hand sides is still successful. We now experiment with solving the extra right-hand side to different levels of accuracy. Table 4.2 shows the relative accuracy (residual norm divided by the initial residual) after correction for the $\sigma = -2$ system when the extra right-hand side system is solved to relative residual tolerances ranging from $1.e-6$ down to $1.e-1$ (the tolerance is checked for termination only at the end of GMRES cycles). The first

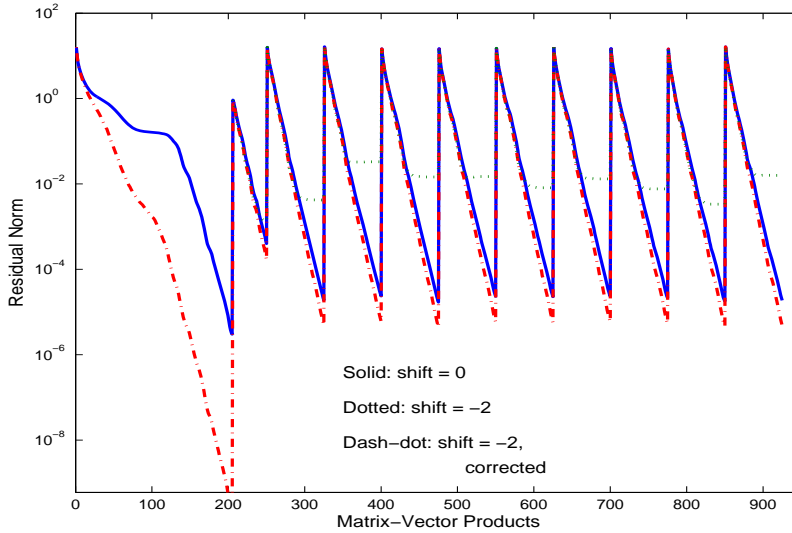


FIG. 4.2. Solution of ten right-hand sides with two shifts.

TABLE 4.2
Effect of solving the extra right-hand side system to different accuracies

desired rtol of 1st and 2nd sys's	rel. accur. of 2nd before correction	1.e-6	1.e-5	1.e-4	1.e-3	1.e-2	1.e-1
1.e-6	1.3e-4	1.5e-7	1.5e-7	1.5e-7	1.5e-7	2.0e-7	1.2e-5
1.e-8	1.1e-5	7.5e-10	7.5e-10	7.5e-10	7.9e-9	3.0e-8	3.1e-6
1.e-10	3.7e-5	4.7e-12	8.5e-12	3.1e-11	3.1e-9	8.5e-9	9.7e-7

and second right-hand side systems are solved to three different relative residual norm tolerances (1.e-6, 1.e-8 and 1.e-10) in the three rows of the table. The conclusion of this experiment is that the extra right-hand side systems do not need to be solved very accurately. For instance, with desired tolerance 1.e-8 for the first and second right-hand sides, the extra right-hand side systems need only to be solved to relative tolerance of 1.e-3 to get relative accuracy of 7.9e-9 for the second right-hand side system.

The next example is probably the key example in the paper. It shows the value of deflating eigenvalues for an important application. QCD problems often have the need of solution of multiple right-hand sides and multiple shifts for each right-hand side. They also have complex spectra with small eigenvalues for the problems of most interest.

Example 6. We look at a Wilson-Dirac matrix from lattice QCD [30]. The matrix is complex and the dimension is 393,216 by 393,216. The value of κ is 0.158 for the base shift. This is approximately $\kappa_{critical}$. The eigenvalues are in the right-half of the complex plane, but partially surround the origin [30]. The right-hand sides are unit vectors associated with particular space-time, Dirac and color coordinates. Often there are a dozen or more right-hand sides associated with each matrix and perhaps seven shifts for each right-hand side. This test was done on the Baylor High Performance Cluster in parallel FORTRAN with calls to LAPACK for solution of the

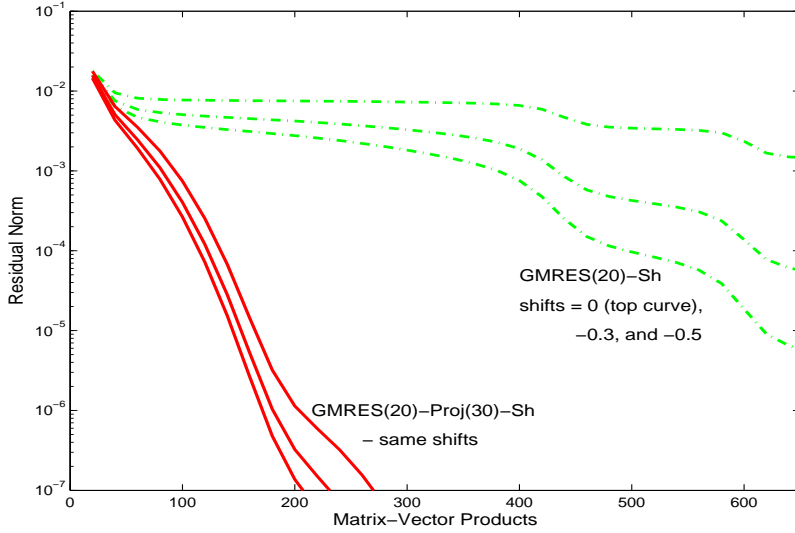


FIG. 4.3. *Solution of the second right-hand side of a QCD matrix.*

eigenvalue problems. We will just show solution of the second right-hand side for three shifts, $\sigma = 0, -0.3, -0.5$. The first right-hand side is solved with GMRES-DR(50,30) to residual tolerance of $1.e-8$ and the extra right-hand side to $1.e-7$. Then for the second right-hand side, GMRES(20)-Proj(30)-Sh uses 30 approximate eigenvectors for the projection in between cycles of GMRES(20)-Sh. See Figure 4.3 for the results. GMRES(20)-Proj(30)-Sh converges in about one-tenth of the iterations needed for GMRES(20)-Sh. To reach residual norm of less than 10^{-7} for the toughest system with shift of zero takes 2680 matrix-vector products for GMRES(20)-Sh and 280 for GMRES(20)-Proj(30)-Sh.

5. Related Right-hand Sides. We need to take advantage of any relationship between the right-hand sides. This section shows that this can easily be done even for the case of multiple shifts.

Assume we have solutions of all shifted systems for right-hand sides 1 through $j - 1$. So we have

$$(A - \sigma_i I)x_i^{irhs} = b^{irhs},$$

for $irhs = 1, \dots, j - 1$. We put these equations together to form

$$(5.1) \quad (A - \sigma_i I)X_i = B,$$

where B is the n by $j - 1$ matrix with columns b^1 through b^{j-1} , and X_i has columns x_i^1 through x_i^{j-1} . We assume there is no initial guess. Applying Minres projection over the subspace spanned by the columns of X_i to the system with the right-hand side j and shift σ_i gives

$$X_i^H (A - \sigma_i I)^H (A - \sigma_i I) X_i d = X_i^H (A - \sigma_i I)^H b^j.$$

With (5.1), this becomes

$$(5.2) \quad B^H B d = B^H b^j.$$

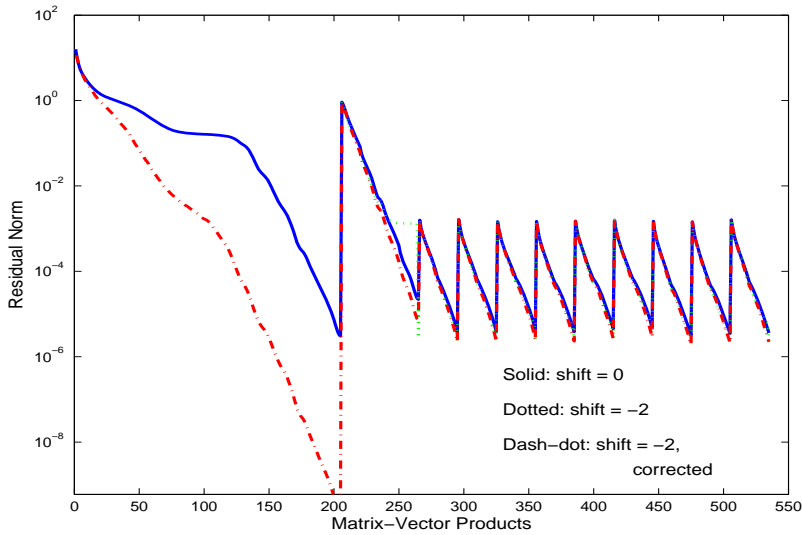


FIG. 5.1. *Solution of ten related right-hand sides.*

Note that the solution of Equation (5.2) is independent of the shift. This makes the residual vectors all the same:

$$r_i = b^j - Bd.$$

The approximate solutions are

$$\tilde{x}_i^j = X_i d.$$

So this approach projects over all of the previous solutions for each shifted system and provides the needed parallel residuals.

Example 7. We repeat the test in Example 4 with 10 right-hand sides, except this time they are related to each other. We define the second and subsequent right-hand sides as $b^j = b^1 + 10^{-4} * u^j$, where u^j is a random vector (both b^1 and u^j have elements distributed normally with mean 0 and variance 1). Before solving the second and subsequent right-hand sides, we project over the previous solutions as just described. The results are in Figure 5.1. Using the close relationship between the right-hand sides allows the number of matrix-vector products to be cut in half for each of the subsequent right-hand sides.

6. Conclusion. Deflating eigenvalues can significantly improve restarted GMRES for matrices with small eigenvalues. This work focuses on deflated GMRES for the case of multiple shifts of the matrix. When using Krylov methods to solve systems of equations with multiple shifts, the goal is to solve all systems with about the same expense as one system. Past work has developed such methods for non-restarted and even restarted Krylov methods. Here this is extended for deflated, restarted GMRES methods, both in the case of a single right-hand side and for multiple right-hand sides. Deflating eigenvalues not only can improve convergence, but it reduces concerns of having shifts near eigenvalues of the base system.

For multiple right-hand sides, there is added expense for solving an extra right-hand side, but this allows all shifted systems to be solved simultaneously for as many right-hand sides as are desired. Also, it is possible to efficiently take advantage of closely related right-hand sides, even in this case of multiple shifts. These approaches can be very beneficial for an important application in lattice QCD physics.

Future work could examine other QCD problems including the overlap fermion problem, which involves solving an inner-outer loop. There may be potential for deflation in both of the loops. Deflation of non-restarted methods such as BiCGStab should also be investigated. Another interesting topic would be solution of systems with changing shifts, such as may occur in model order reduction [22].

Acknowledgments. We thank the referees for many corrections and helpful suggestions.

REFERENCES

- [1] G. Arnold, N. Cundy, J. van den Eshof, A. Frommer, S. Krieg, Th. Lippert, and K. Sch fer. Numerical methods for the QCD overlap operator: II. Sign-function and error bounds. 2003.
- [2] J. Baglama, D. Calvetti, G. H. Golub, and L. Reichel. Adaptively preconditioned GMRES algorithms. *SIAM J. Sci. Comput.*, 20:243–269, 1998.
- [3] K. Burrage and J. Erhel. On the performance of various adaptive preconditioned GMRES strategies. *Num. Lin. Alg. with Appl.*, 5:101–121, 1998.
- [4] C. Le Calvez and B. Molina. Implicitly restarted and deflated GMRES. *Numer. Algo.*, 21:261–285, 1999.
- [5] T. F. Chan and W. Wan. Analysis of projection methods for solving linear systems with multiple right-hand sides. *SIAM J. Sci. Comput.*, 18:1698–1721, 1997.
- [6] A. Chapman and Y. Saad. Deflated and augmented Krylov subspace techniques. *Num. Lin. Alg. with Appl.*, 4:43–66, 1997.
- [7] D. Darnell, R. B. Morgan, and W. Wilcox. Deflation of eigenvalues for iterative methods in lattice QCD. *Nucl. Phys. B (Proc. Suppl.)*, 129:856–858, 2004.
- [8] B. N. Datta and Y. Saad. Sylvester-like observer matrix equations, and an associated algorithm for partial spectrum assignment. *Linear Algebra Appl.*, 154-156:225–244, 1991.
- [9] P. de Forcrand. Progress on lattice QCD algorithms. *Nucl. Phys. B (Proc. Suppl.)*, 47:228–235, 1996.
- [10] E. De Sturler. Truncation strategies for optimal Krylov subspace methods. *SIAM J. Numer. Anal.*, 36:864–889, 1999.
- [11] S. J. Dong, F. X. Lee, K. F. Liu, and J. B. Zhang. Chiral symmetry, quark mass, and scaling of the overlap fermions. *Phys. Rev. Lett.*, 85:5051–5054, 2000.
- [12] R. G. Edwards, U. M. Heller, and R. Narayanan. Study of chiral symmetry in quenched QCD using the overlap Dirac operator. *Phys. Rev. D*, 59:0945101–0945108, 1999.
- [13] J. Erhel, K. Burrage, and B. Pohl. Restarted GMRES preconditioned by deflation. *J. Comput. Appl. Math.*, 69:303–318, 1996.
- [14] A. Feriani, F. Perotti, and V. Simoncini. Iterative system solvers for the frequency analysis of linear mechanical systems. *Comput. Meth. Appl. Mech. and Engrg.*, 190:1719–1739, 2000.
- [15] R. W. Freund. On conjugate gradient type methods and polynomial preconditioners for a class of complex non-Hermitian matrices. *Numer. Math.*, 57:285–312, 1990.
- [16] R. W. Freund. Quasi-kernel polynomials and their use in non-Hermitian matrix iterations. *J. Comput. Appl. Math.*, 43:135–158, 1992.
- [17] R. W. Freund. Solution of shifted linear systems by quasi-minimal residual iterations. In L. Reichel, A. Ruttan, and R. S. Varga, editors, *Numerical Linear Algebra*, pages 101–121. W. de Gruyter, 1993.
- [18] R. W. Freund and M. Malhotra. A block QMR algorithm for non-Hermitian linear systems with multiple right-hand sides. *Linear Algebra Appl.*, 254:119–157, 1997.
- [19] A. Frommer. BiCGStab(l) for families of shifted linear systems. *Computing*, 70:87–109, 2003.
- [20] A. Frommer and U. Glässner. Restarted GMRES for shifted linear systems. *SIAM J. Sci. Comput.*, 19:15–26, 1998.
- [21] G. D. Gu, J. J. Zhang, and Z. W. Li. Restarted GMRES augmented with eigenvectors for shifted linear systems. *Inter. J. Comput. Math.*, 80:1037–1047, 2003.

- [22] S. Gugercin, A.C. Antoulas, and C.A. Beattie. Rational Krylov methods for optimal H2 model reduction. 2006.
- [23] S. A. Kharchenko and A. Y. Yeremin. Eigenvalue translation based preconditioners for the GMRES(k) method. *Num. Lin. Alg. with Appl.*, 2:51–77, 1995.
- [24] R. B. Morgan. Computing interior eigenvalues of large matrices. *Linear Algebra Appl.*, 154-156:289–309, 1991.
- [25] R. B. Morgan. A restarted GMRES method augmented with eigenvectors. *SIAM J. Matrix Anal. Appl.*, 16:1154–1171, 1995.
- [26] R. B. Morgan. On restarting the Arnoldi method for large nonsymmetric eigenvalue problems. *Math. Comp.*, 65:1213–1230, 1996.
- [27] R. B. Morgan. Implicitly restarted GMRES and Arnoldi methods for nonsymmetric systems of equations. *SIAM J. Matrix Anal. Appl.*, 21:1112–1135, 2000.
- [28] R. B. Morgan. GMRES with deflated restarting. *SIAM J. Sci. Comput.*, 24:20–37, 2002.
- [29] R. B. Morgan. Restarted block-GMRES with deflation of eigenvalues. *Appl. Numer. Math.*, 54:222–236, 2005.
- [30] R. B. Morgan and W. Wilcox. Deflation of eigenvalues for GMRES in lattice QCD. *Nucl. Phys. B (Proc. Suppl.)*, 106:1067–1069, 2002.
- [31] R. B. Morgan and W. Wilcox. Deflated iterative methods for linear equations with multiple right-hand sides. arXiv:math-ph/0405053v2, 2004.
- [32] R. B. Morgan and M. Zeng. Harmonic projection methods for large non-symmetric eigenvalue problems. *Numer. Linear Algebra Appl.*, 5:33–55, 1998.
- [33] R. B. Morgan and M. Zeng. A harmonic restarted Arnoldi for calculating eigenvalues and determining multiplicity. *Linear Algebra Appl.*, 415:96–113, 2006.
- [34] R. Narayanan and H. Neuberger. An alternative to domain wall fermions. *Phys. Rev. D*, 62:074504, 2000.
- [35] H. Neff, N. Eicker, Th. Lippert, J. W. Negele, and K. Schilling. On the low fermionic eigenmode dominance in QCD on the lattice. *Phys. Rev. D*, 64:114509–1 – 114509–12, 2001.
- [36] C. C. Paige, B. N. Parlett, and H. A. van der Vorst. Approximate solutions and eigenvalue bounds from Krylov subspaces. *Num. Lin. Alg. with Appl.*, 2:115–133, 1995.
- [37] Michael L. Parks, Eric de Sturler, Greg Mackey, Duane D. Johnson, and Spandan Maiti. Recycling Krylov subspaces for sequences of linear systems. *SIAM J. Sci. Comput.*, 28:1651–1674, 2006.
- [38] Y. Saad. *Iterative Methods for Sparse Linear Systems*. PWS Publishing, Boston, MA, 1996.
- [39] Y. Saad. Analysis of augmented Krylov subspace techniques. *SIAM J. Matrix Anal. Appl.*, 18:435–449, 1997.
- [40] Y. Saad and M. H. Schultz. GMRES: a generalized minimum residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Statist. Comput.*, 7:856–869, 1986.
- [41] V. Simoncini. Restarted full orthogonalization method for shifted linear systems. *BIT*, 43:459–466, 2003.
- [42] V. Simoncini and E. Gallopoulos. An iterative method for nonsymmetric systems with multiple right-hand sides. *SIAM J. Sci. Comput.*, 16:917–933, 1995.
- [43] V. Simoncini and E. Gallopoulos. A hybrid block GMRES method for nonsymmetric systems with multiple right-hand sides. *J. Comp. and Applied Math.*, 66:457–469, 1996.
- [44] D. C. Sorensen. Implicit application of polynomial filters in a k -step Arnoldi method. *SIAM J. Matrix Anal. Appl.*, 13:357–385, 1992.
- [45] G. W. Stewart. A Krylov–Schur algorithm for large eigenproblems. *SIAM J. Matrix Anal. Appl.*, 23:601 – 614, 2001.
- [46] K. Wu and H. Simon. Thick-restart Lanczos method for symmetric eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 22:602 – 616, 2000.
- [47] D. M. Young and B. R. Vona. On the use of rational iterative methods for solving large sparse linear systems. *Appl. Numer. Math.*, 10:261 – 278, 1992.