

## PROJECTED KRYLOV METHODS FOR SADDLE-POINT SYSTEMS\*

NICK GOULD<sup>†</sup>, DOMINIQUE ORBAN<sup>‡</sup>, AND TYRONE REES<sup>†</sup>

**Abstract.** Projected Krylov methods are full-space formulations of Krylov methods that take place in a nullspace. Provided projections into the nullspace can be computed accurately, those methods only require products between an operator and vectors lying in the nullspace. We provide systematic principles for obtaining the projected form of any well-defined Krylov method. Projected Krylov methods are mathematically equivalent to constraint-preconditioned Krylov methods provided the initial guess is well chosen, but require less memory. As a consequence, there are situations where certain known methods such as MINRES and SYMMLQ are well defined in the presence of an indefinite preconditioner.

**Key words.** saddle-point system, projected Krylov method, constraint preconditioner

**AMS subject classifications.** 15A06, 65F08, 65F10, 65F50, 65H10, 90C20

**DOI.** 10.1137/130916394

**1. Introduction.** We consider the solution of the saddle-point problem

$$(1.1) \quad \begin{bmatrix} \mathbf{Q} & \mathbf{A}^T \\ \mathbf{A} & \end{bmatrix} \begin{bmatrix} \mathbf{x}_* + \mathbf{x}_F \\ \mathbf{y}_* \end{bmatrix} = \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix}$$

for some fixed  $\mathbf{x}_F$ , where all data are real and  $\mathbf{Q}$  may be unsymmetric. Such saddle-point systems arise throughout computational science (Benzi, Golub, and Liesen, 2005). Typical applications arise in conservative fluid flow, structural engineering, and constrained optimization. The scale of the problem often precludes a direct factorization of the matrix  $\mathbf{K}_Q$  in (1.1). Certain applications give rise to a symmetric saddle-point system, including optimization and the study of laminar fluid flow. A common approach in such cases is to employ a Krylov method for symmetric indefinite systems, e.g., MINRES or SYMMLQ (Paige and Saunders, 1975), combined with an appropriate preconditioner (Elman, Silvester, and Wathen, 2005).

In this paper we build upon the work of Gould, Hribar, and Nocedal (2001) and highlight families of iterative methods working implicitly in the nullspace of  $\mathbf{A}$ , requiring only operator-vector products with  $\mathbf{Q}$  (and possibly its transpose) as well as one or more projections of a vector into the nullspace of  $\mathbf{A}$  per iteration. We require that  $\mathbf{Q}$  be available as an operator, and need accurate projections into the nullspace of  $\mathbf{A}$ . Gould, Hribar, and Nocedal (2001) describe how to compute sufficiently accurate projections at moderate cost by way of iterative refinement or residual updates.

Our main contribution is to provide systematic principles to derive a projected variant from any well-defined Krylov method. This is made possible by working at the level of the basis construction processes upon which those methods are built. An additional contribution is to provide equivalence relations in the general case between

---

\*Received by the editors April 10, 2013; accepted for publication (in revised form) by M. Benzi August 20, 2014; published electronically November 4, 2014.

<http://www.siam.org/journals/simax/35-4/91639.html>

<sup>†</sup>Scientific Computing Department, Rutherford Appleton Laboratory, Chilton, Oxfordshire OX11 0QX, England (nick.gould@stfc.ac.uk, tyrone.rees@stfc.ac.uk). The research of these authors was supported by EPSRC grant EP/I013067/1.

<sup>‡</sup>GERAD and Department of Mathematics and Industrial Engineering, École Polytechnique, Montréal, QC, Canada (dominique.orban@gerad.ca). The research of this author was partially supported by an NSERC Discovery grant.

projected Krylov methods and traditional Krylov methods applied directly to (1.1) with a so-called *constraint* preconditioner.

As we mention in section 4, the principles highlighted below are applicable beyond the strict context of Krylov methods. In particular, they may be applied to derive projected variants of iterative methods based on Krylov-type basis-generation processes. An example is given by Orban (2014), who derives a projected variant of the Golub and Kahan (1965) process and applies it to equality-constrained least-squares problems.

Our approach involves a sequence of key steps. We first reduce the saddle-point problem to an equivalent one in the nullspace of  $\mathbf{A}$  and then transform (precondition) the result. We next apply an appropriate Krylov method to the resulting preconditioned, reduced system. The effects of the preconditioner on the method are then considered in the unpreconditioned reduced space, and finally the iteration is moved back from the nullspace via a constraint preconditioner into its original full-space setting. Our framework is closely related to the *nullspace method*—see (Benzi, Golub, and Liesen, 2005, section 6) and references therein—but differs from it in that we only use the data of (1.1). Similarly to Gould, Hribar, and Nocedal (2001), we do not require that a basis for  $\text{Null}(\mathbf{A})$  be computed.

An interesting consequence of our framework is to establish that well-known methods such as MINRES and SYMMLQ applied to (1.1) may be well defined in the presence of an indefinite preconditioner. Specifically, MINRES minimizes the residual in the seminorm defined by the constraint preconditioner which, due to the structure of the residuals, results in a well-defined method—we give more detail in section 3. Similar considerations hold for SYMMLQ, and the equivalent result was shown by Rozložník and Simoncini (2002) to be true for definite systems. This is at variance with the commonly issued warnings (e.g., Elman, Silvester, and Wathen, 2005, p. 287; Greenbaum, 1997, p. 121; and van der Vorst, 2003, p. 85) that MINRES and SYMMLQ require definite preconditioners.

In section 2, we examine a number of well-known Krylov-subspace methods for saddle-point systems. Specifically, in section 2.1, we show how methods may be constructed to reflect the saddle-point structure; in section 2.2, we detail the bases used and how standard methods appear when moved into the subspace defined by such structure; and in section 2.3 we return these methods back into the original space. In section 2.4 we show that our construction is equivalent to applying standard methods with a constraint preconditioner. We illustrate this equivalence on the MINRES method in section 3 by showing that a positive definite preconditioner is not required when applying the method to saddle-point systems. We briefly consider other variants in section 4, and conclude in section 5.

**1.1. Related research.** In optimization contexts, where  $\mathbf{Q}$  is symmetric, (1.1) may be interpreted as the first-order optimality conditions of the quadratic program

$$(1.2) \quad \underset{\mathbf{x} + \mathbf{x}_F}{\text{minimize}} \quad -\mathbf{a}^\top(\mathbf{x} + \mathbf{x}_F) + \frac{1}{2}(\mathbf{x} + \mathbf{x}_F)^\top \mathbf{Q}(\mathbf{x} + \mathbf{x}_F) \quad \text{subject to} \quad \mathbf{A}(\mathbf{x} + \mathbf{x}_F) = \mathbf{b},$$

where  $\mathbf{y}$  are the Lagrange multipliers associated with the equality constraints. It is well known (see, e.g., Gould, 1985), that (1.2) possesses a unique solution if and only if  $\mathbf{Q}$  is positive definite on the nullspace of  $\mathbf{A}$ . Based on this, Gould, Hribar, and Nocedal (2001) devised the projected conjugate gradient method, a variant of the standard conjugate gradient algorithm applied to  $\mathbf{Q}$  and restricted to exploration of the nullspace of  $\mathbf{A}$ —see also Polyak (1969), Coleman (1994), Lukšan and Vlček (1998), and Perugia

and Simoncini (2000). Benzi, Golub, and Liesen (2005, section 6) describe and provide numerous references on the nullspace method, which requires the computation of a basis for  $\text{Null}(\mathbf{A})$ . Numerical challenges quickly accumulate as desirable properties for this basis, such as sparsity and good conditioning, typically come at high expense. As it turns out, Gould, Hribar, and Nocedal (2001) show that it is possible to avoid computing a basis for the nullspace of  $\mathbf{A}$  altogether and formulate the entire algorithm in terms of projections into this nullspace and operator-vector products with  $\mathbf{Q}$ . One way to compute projections efficiently when  $\mathbf{A}$  is available as an explicit matrix is to perform a one-time factorization of a symmetric indefinite matrix of the form (1.1) where  $\mathbf{Q}$  is replaced by a simpler operator such as (but not restricted to) the identity. In many applications, performing this factorization is realistic and cost effective; for example, one may choose the (1,1) block so that a convenient Schillers' factorization (Dollar and Wathen, 2006) may be employed.

Orban (2008) applies the same principles as those provided by Gould, Hribar, and Nocedal (2001) to specific Krylov methods for unsymmetric systems with the solution of fluid flow problems in mind. He focuses on methods not requiring products with the transpose operator.

**1.2. Terminology and notation.** We refer to the quantities  $\mathbf{Q}$ ,  $\mathbf{A}$ ,  $\mathbf{x}$ , etc., of (1.1) as *full-space* quantities. Throughout the paper, they are typeset in italicized bold Roman font. Whenever we work explicitly in the nullspace of  $\mathbf{A}$ , we refer to corresponding quantities as *reduced-space* quantities and typeset them in italicized Roman light face font, e.g.,  $Q$ ,  $A$ ,  $x$ . We precondition reduced-space quantities and refer to them as *preconditioned-space* quantities. They are typeset in an italicized bold sans-serif font, e.g.,  $\mathbf{Q}$ ,  $\mathbf{A}$ ,  $\mathbf{x}$ . The Euclidean norm and its associated inner product are denoted  $\|\cdot\|$  and  $\langle \cdot, \cdot \rangle$ , respectively, throughout.

**2. Saddle-point problems.** Throughout the paper, our working assumption is the following.

*Assumption 2.1.* The coefficient matrix of (1.1) is nonsingular.

Benzi, Golub, and Liesen (2005, Theorem 3.4) provide necessary and sufficient conditions for Assumption 2.1 to hold. In particular, if  $\mathbf{H} = \frac{1}{2}(\mathbf{Q} + \mathbf{Q}^T)$  denotes the symmetric part of  $\mathbf{Q}$ , and if we assume that  $\mathbf{H}$  is positive semidefinite and  $\mathbf{A}$  has full row rank, then

- (1.1) is nonsingular if  $\text{Null}(\mathbf{H}) \cap \text{Null}(\mathbf{A}) = \{\mathbf{0}\}$ , and
- $\text{Null}(\mathbf{Q}) \cap \text{Null}(\mathbf{A}) = \{\mathbf{0}\}$  if (1.1) is nonsingular,

but the reverse implications do not hold in general. The above conditions are sufficiently general to encompass numerous applications of interest; we refer the interested reader to (Benzi, Golub, and Liesen, 2005) for more details. At least in theory, it is possible to ensure that  $\mathbf{A}$  has full row rank by way of preprocessing. In some cases, regularization is preferred and consists in giving a nonzero value to the (2, 2) block of (1.1). In section 4, we explain how the methods presented in the next sections apply to such a regularized system.

Suppose that  $\mathbf{x}_F$  satisfies  $\mathbf{A}\mathbf{x}_F = \mathbf{b}$ ; as we will see, finding such a value is easy in the framework we develop. Substituting into (1.1),  $\mathbf{x}_*$  and  $\mathbf{y}_*$  satisfy

$$(2.1) \quad \begin{bmatrix} \mathbf{Q} & \mathbf{A}^T \\ \mathbf{A} & \end{bmatrix} \begin{bmatrix} \mathbf{x}_* \\ \mathbf{y}_* \end{bmatrix} = \begin{bmatrix} \mathbf{c} \\ \mathbf{0} \end{bmatrix},$$

where  $\mathbf{c} := \mathbf{a} - \mathbf{Q}\mathbf{x}_F$ . Our aim is thus to solve (2.1).

Let  $\mathbf{Z}$  be any full-rank matrix whose columns span the nullspace of  $\mathbf{A}$ , i.e., such that

$$(2.2) \quad \mathbf{AZ} = \mathbf{0}.$$

Then necessarily from the second block of (2.1),  $\mathbf{x}_* = \mathbf{Z}x$  for some  $x$ , in which case the first block gives  $\mathbf{QZ}x + \mathbf{A}^\top \mathbf{y}_* = \mathbf{c}$  and, hence, from (2.2),

$$(2.3) \quad \mathbf{Q}x_* = \mathbf{c}, \quad \text{where } \mathbf{Q} := \mathbf{Z}^\top \mathbf{QZ} \quad \text{and} \quad \mathbf{c} := \mathbf{Z}^\top \mathbf{c}.$$

Notice that if we define

$$\mathbf{r} := \mathbf{c} - \mathbf{Q}\mathbf{x} - \mathbf{A}^\top \mathbf{y} \quad \text{and} \quad r := \mathbf{c} - \mathbf{Q}x,$$

it follows immediately from (2.2) that residuals are transformed according to

$$(2.4) \quad r = \mathbf{Z}^\top \mathbf{r}$$

regardless of the value of  $\mathbf{y}$ .

Let  $G$  be a symmetric positive-definite approximation to  $\mathbf{Q}$ , and suppose that we may write  $G = LL^\top$  as necessary. We will consider the central-preconditioned variant

$$(2.5) \quad L^{-1}\mathbf{Q}L^{-\top}(L^\top x_*) = L^{-1}\mathbf{c}$$

of (2.3). If we define

$$(2.6) \quad \mathbf{Q} := L^{-1}\mathbf{Q}L^{-\top}, \quad \mathbf{c} := L^{-1}\mathbf{c}, \quad \text{and} \quad \mathbf{x}_* = L^{-\top}x_*,$$

the system (2.5) may be expressed compactly as

$$\mathbf{Q}\mathbf{x}_* = \mathbf{c}.$$

Here, residuals are transformed according to

$$r = \mathbf{c} - \mathbf{Q}\mathbf{x} = L^{-1}(\mathbf{c} - \mathbf{Q}\mathbf{x}) = L^{-1}r.$$

**2.1. Krylov-subspace methods.** The  $k$ th approximations  $\mathbf{x}_k \approx \mathbf{x}_*$  generated by Krylov methods lie in either of the  $k$ -dimensional Krylov spaces

$$(2.7) \quad \mathcal{K}_k := \text{Span} \{ \mathbf{Q}^i \mathbf{c} \}_{i=0}^{k-1} \quad \text{or} \quad \mathcal{K}_k^\top := \text{Span} \{ (\mathbf{Q}^\top)^i \mathbf{d} \}_{i=0}^{k-1}$$

for specified  $\mathbf{d}$  for which  $\langle \mathbf{c}, \mathbf{d} \rangle \neq 0$ . Let  $\mathbf{V}_k$  be a matrix whose columns form a basis of  $\mathcal{K}_k$ . The choice of  $\mathbf{x}_k$  is characterized by the basis of  $\mathcal{K}_k$  that is chosen and an optimality condition. (*Ritz-)*Galerkin methods seek solutions of the form  $\mathbf{x}_k := \mathbf{V}_k \mathbf{z}$  and impose that the residual  $\mathbf{r}_k := \mathbf{c} - \mathbf{Q}\mathbf{x}_k$  be orthogonal to  $\mathcal{K}_k$ . *Minimum-residual* methods also seek  $\mathbf{x}_k := \mathbf{V}_k \mathbf{z}$ , where  $\mathbf{z} \in \arg \min_{\mathbf{z}} \|\mathbf{Q}\mathbf{V}_k \mathbf{z} - \mathbf{c}\|$  in some specified norm. *Minimum-error* methods seek approximations  $\mathbf{x}_k := \mathbf{Q}^\top \mathbf{V}_k \mathbf{z}$ , where  $\mathbf{z} \in \arg \min_{\mathbf{z}} \|\mathbf{Q}^\top \mathbf{V}_k \mathbf{z} - \mathbf{x}_*\|$ .

The Arnoldi (1951) process may be used to generate orthogonal basis matrices  $\mathbf{V}_k^\circ$  and  $\mathbf{W}_k^\circ$  of  $\mathcal{K}_k$  and  $\mathcal{K}_k^\top$ , respectively. When  $\mathbf{Q}$  is symmetric, the Arnoldi process simplifies into the symmetric Lanczos (1950) process. An alternative is to use a biorthogonal pair of basis matrices  $\mathbf{V}_k$  and  $\mathbf{W}_k$  for  $\mathcal{K}_k$  and  $\mathcal{K}_k^\top$ . This basis-generation process is known as the Lanczos (1950) biorthogonalization process. It is a riskier alternative and may break down, but when  $\mathbf{Q}$  is symmetric, breakdown will not occur if  $\mathbf{c} = \mathbf{d}$  (van der Vorst, 2003, section 7.1). A good general overview is provided by Saad (2003).

Each of the possibilities depends entirely on its basis matrix  $V_k^\circ$  or  $V_k$  and the approximation approach. We concentrate on general Krylov methods for which

$$x_k = V_k^\circ z_k \quad \text{or} \quad x_k = V_k z_k,$$

or, in terms of the unpreconditioned variables,

$$x_k = V_k^\circ z_k \quad \text{or} \quad x_k = V_k z_k,$$

where

$$(2.8) \quad V_k^\circ = L^{-T} V_k^\circ \quad \text{and} \quad V_k = L^{-T} V_k$$

and  $z_k$  is computed by any means from appropriate data  $Z_k$ ; we shall formally denote the algorithm used to compute  $z_k$  as

$$(2.9) \quad z_k = \Theta_k(Z_k),$$

as appropriate, for some vector valued function  $\Theta_k$  of dimension  $k$ .

Note that the vectors in  $V_k^\circ$  are not orthogonal in the  $\ell_2$  norm, but will be in the  $G^{-1}$  norm. We inherit the notation from the orthogonal matrix  $V_k^\circ$  with this understanding. We denote the columns of  $V_k^\circ$  by  $v_i^\circ, i = 0, \dots, k - 1$ .

**2.2. Computing suitable bases.** The other main ingredient in our development is to recall how the bases we mentioned above are computed. Beginning with standard methods for  $Q$  as given by (2.6) we derive the corresponding iteration for the data  $Q$  and its preconditioner  $G$ . In the next few sections, we focus on the Arnoldi process, but similar principles may be applied to other basis-generation processes. We leave until section 2.3 a detailed discussion on how these computed bases may be applied in the full space to solve (2.1).

For a given orthogonalization method, in order to formulate corresponding algorithm involving unpreconditioned quantities we apply changes of variables suggested by (2.6) and (2.8). All transformations in the next few sections will follow the same principle.

PRINCIPLE 2.1.

1. *Basis vectors transform according to  $v_k^\circ = L^{-T} v_k^\circ$ . Because of (2.6) and (2.7), each generated column  $v_k^\circ$  of  $V_k^\circ$  in preconditioned space has the form  $v_k^\circ = L^{-1}u$  for some vector  $u$ . Thus, we obtain  $v_k^\circ$  by solving the preconditioning system  $Gv_k^\circ = u$  for  $v_k^\circ$ .*
2. *Inner products of the form  $\langle v_k^\circ, v_i^\circ \rangle$  become  $\langle u, v_i^\circ \rangle$ , again because of (2.6), where  $u$  is the vector defined as in the first part of the principle.*

The Arnoldi process applied to (2.6) ensures that

$$(2.10) \quad QV_k = V_{k+1}H_{k+1,k}$$

$$(2.11) \quad = V_k H_k + h_{k+1,k} v_{k+1} e_k^T,$$

where  $H_k$  is  $k$ -by- $k$  upper Hessenberg and  $H_{k+1,k}$  is  $H_k$  with the extra row  $h_{k+1,k} e_k^T$ . Applying Principle 2.1 to this we obtain Algorithm 2.1 in which everything is expressed in the reduced (unpreconditioned) space. Note also that there is now no explicit mention of  $L$ , but rather that solves with  $G$ , using any appropriate mechanism, are required.

---

ALGORITHM 2.1. PRECONDITIONED ARNOLDI PROCESS FOR  $V_k^\circ$  AND  $H_{k,k-1}$ .

---

**Require:**  $Q$ ,  $G = G^\top \succ 0$ ,  $c$  and  $x_0$

```

1:  $u = c - Qx_0$ 
2: Solve  $Gv_1^\circ = u$  for  $v_1^\circ$  // Initial Krylov vector
3:  $h_{1,0} = \sqrt{\langle u, v_1^\circ \rangle}$  // Initial preconditioned residual norm
4: if  $h_{1,0} \neq 0$  then
5:    $v_1^\circ = v_1^\circ / h_{1,0}$ 
6:  $k = 1$ 
7: while  $h_{k,k-1} \neq 0$  do
8:    $u = Qv_k^\circ$  and Solve  $Gv_{k+1}^\circ = u$  // Compute next Krylov vector
9:   for  $i = 1, \dots, k$  do // Modified Gram-Schmidt
10:     $h_{i,k} = \langle u, v_i^\circ \rangle$ 
11:     $v_{k+1}^\circ = v_{k+1}^\circ - h_{i,k}v_i^\circ$ 
12:    $h_{k+1,k} = \sqrt{\langle u, v_{k+1}^\circ \rangle}$ 
13:   if  $h_{k+1,k} \neq 0$  then
14:     $v_{k+1}^\circ = v_{k+1}^\circ / h_{k+1,k}$ 
15:    $k = k + 1$ 

```

---

Principle 2.1 is directly applicable to other basis-generation processes. For instance, it may be applied verbatim to the Lanczos (1950) biorthogonalization process to obtain a corresponding process formulated in terms of reduced-space quantities and the preconditioner  $G$ . It may also be applied to the variant given by Chan, de Pillis, and van der Vorst (1998) and Freund, Gutknecht, and Nachtigal (1993), or to the transpose-free variant (Brezinski and Redivo-Zaglia, 1998; Chan, de Pillis, and van der Vorst, 1998).

**2.3. Iteration in the full space.** We are now in a position to describe how the Krylov methods we have considered in the nullspace of  $\mathbf{A}$  may actually be applied in the original (full) space. This is achieved via a systematic principle that may be applied to any basis-generation process. The outcome of the application of the principle was previously known in special cases, such as the case where  $\mathbf{Q}$  is symmetric and positive definite, e.g., (Gould, Hribar, and Nocedal, 2001). Our description is general in the sense that it applies to any valid Krylov method. Recall that in the full space we have

$$\mathbf{x}_* = \mathbf{Z}x_*, \quad \mathbf{Q} = \mathbf{Z}^\top \mathbf{Q} \mathbf{Z}, \quad \text{and} \quad \mathbf{c} = \mathbf{Z}^\top \mathbf{c},$$

in which case we have

$$\mathbf{x}_k = \mathbf{Z}x_k = \mathbf{Z}V_k^\circ z_k \quad \text{or} \quad \mathbf{x}_k = \mathbf{Z}x_k = \mathbf{Z}V_k z_k.$$

Consider also a preconditioner of the form

$$G := \mathbf{Z}^\top \mathbf{G} \mathbf{Z},$$

where  $\mathbf{G}$  satisfies the following requirement.

*Assumption 2.2.* The matrix  $\mathbf{G}$  is symmetric and positive definite on  $\text{Null}(\mathbf{A})$ .

Such a requirement on  $\mathbf{G}$  clearly implies that  $G$  is symmetric, positive definite. But while choosing  $\mathbf{G}$  itself to be positive definite will suffice, this is far from necessary and indeed undesirable if, as is common for constrained optimization,  $\mathbf{Q}$  is indefinite.

Consider iterates generated according to  $\mathbf{x}_k = \mathbf{Z}x_k = \mathbf{Z}V_k^\circ z_k$  by the Arnoldi process. We now apply the following principle to express Algorithm 2.1 in terms of full-space quantities.

PRINCIPLE 2.2.

1. Basis vectors transform according to  $\mathbf{v}_k^\circ = \mathbf{Z}\mathbf{v}_k^\circ$ . Because of (2.3), each assignment of  $\mathbf{v}_k^\circ$  in the reduced space has the form  $(\mathbf{Z}^\top \mathbf{G} \mathbf{Z})^{-1} \mathbf{u}$  for some  $\mathbf{u}$  of the form  $\mathbf{u} = \mathbf{Z}^\top \mathbf{u}$ . Thus, we obtain  $\mathbf{v}_k^\circ = \mathbf{Z}(\mathbf{Z}^\top \mathbf{G} \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{u}$ .
2. Inner products of the form  $\langle \mathbf{u}, \mathbf{v}_k^\circ \rangle$  become  $\langle \mathbf{Z}^\top \mathbf{u}, (\mathbf{Z}^\top \mathbf{G} \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{u} \rangle = \langle \mathbf{u}, \mathbf{v}_k^\circ \rangle$ .

At first sight, Principle 2.2 appears to suggest that full-space algorithms depend explicitly on  $\mathbf{Z}$ . However our choice of  $\mathbf{G}$  ensures that the crucial operator

$$(2.12) \quad \mathbf{P}_G := \mathbf{Z}(\mathbf{Z}^\top \mathbf{G} \mathbf{Z})^{-1} \mathbf{Z}^\top$$

has properties akin to those of an orthogonal projector into the nullspace of  $\mathbf{A}$ . More precisely,  $\mathbf{P}_G \mathbf{G}$  is an oblique projector into  $\text{Null}(\mathbf{A})$ .

We summarize a few immediate properties of  $\mathbf{P}_G$  in the following result.

THEOREM 2.1. Let  $\mathbf{P}_G$  be defined as in (2.12) where  $\mathbf{G}$  satisfies Assumption 2.2.

Then

1.  $\mathbf{P}_G \mathbf{G} \mathbf{P}_G = \mathbf{P}_G$  and  $(\mathbf{P}_G \mathbf{G})^2 = \mathbf{P}_G \mathbf{G}$ ,
2.  $\mathbf{P}_G \mathbf{x} = \mathbf{0}$  for all  $\mathbf{x} \in \text{Range}(\mathbf{A}^\top)$ ,
3.  $\mathbf{P}_G \mathbf{G} \mathbf{x} = \mathbf{x}$  for all  $\mathbf{x} \in \text{Null}(\mathbf{A})$ ,
4.  $\mathbf{Z}^\top \mathbf{G} \mathbf{P}_G = \mathbf{Z}^\top$  and  $\mathbf{Z}^\top \mathbf{G}(\mathbf{I} - \mathbf{P}_G \mathbf{G}) = \mathbf{0}$ ,
5.  $\mathbf{P}_G \mathbf{G} \mathbf{Z} = \mathbf{Z}$ .

A consequence of the above is that finding  $\mathbf{Z}$  is unnecessary as it is easy to show (Gould, Hribar, and Nocedal, 2001) that  $\mathbf{v}_k^\circ = \mathbf{P}_G \mathbf{u}$  may instead be computed by solving the *constraint-preconditioned* (Keller, Gould, and Wathen, 2000) saddle-point system

$$(2.13) \quad \begin{bmatrix} \mathbf{G} & \mathbf{A}^\top \\ \mathbf{A} & \end{bmatrix} \begin{bmatrix} \mathbf{v}_k^\circ \\ \mathbf{y}_k \end{bmatrix} = \begin{bmatrix} \mathbf{u} \\ \mathbf{0} \end{bmatrix}.$$

For the coefficient matrix  $\mathbf{K}_G$  of (2.13) to be a constraint preconditioner, all that is needed is that  $\mathbf{G}$  satisfies Assumption 2.2 or, equivalently, that  $\mathbf{K}_G$  be nonsingular and have precisely as many negative eigenvalues as  $\mathbf{A}$  has linearly independent rows. An instructive way to summarize the difference between projected and constraint-preconditioned Krylov methods is that projected methods use the preconditioner

$$\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{A} & \end{bmatrix} \begin{bmatrix} \mathbf{G} & \mathbf{A}^\top \\ \mathbf{A} & \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix}.$$

This expression of the projected preconditioner effectively extracts the leading block of the inverse of the constraint preconditioner, which is indeed a projector into  $\text{Null}(\mathbf{A})$ .

Typically, constraint preconditioners are obtained and used by factorizing—either explicitly Keller, Gould, and Wathen (2000) or implicitly Dollar et al. (2006)— $\mathbf{K}_G$ ; this has the further advantage that the required initial value  $\mathbf{x}_F$  may be found by solving

$$(2.14) \quad \begin{bmatrix} \mathbf{G} & \mathbf{A}^\top \\ \mathbf{A} & \end{bmatrix} \begin{bmatrix} \mathbf{x}_F \\ \mathbf{y}_0 \end{bmatrix} = \begin{bmatrix} \mathbf{g} \\ \mathbf{b} \end{bmatrix}$$

for arbitrary  $\mathbf{g}$ —usually,  $\mathbf{g} = \mathbf{0}$  or  $\mathbf{g} = \mathbf{a}$ .

Applying Principle 2.2 to Algorithm 2.1, we derive the class of Krylov methods for (2.1) described by Algorithm 2.2 below. We note that although formally we may express methods using Algorithm 2.2, an actual implementation will usually be more

---

ALGORITHM 2.2. ARNOLDI-BASED PROJECTED KRYLOV METHODS FOR (2.1).

---

**Require:**  $Q$ ,  $P_G$ ,  $c$  and  $x_0$

```

1:  $u = c - Qx_0$ 
2: Compute  $v_1 = P_G u$  // Initial Krylov vector
3:  $h_{1,0} = \sqrt{\langle v_1, u \rangle}$  // Initial preconditioned residual norm
4: if  $h_{1,0} \neq 0$  then
5:    $v_1 = v_1 / h_{1,0}$ 
6:  $x_1 = V_1^\circ \Theta_1(H_{1,0}, h_{1,0})$  // Initial solution estimate
7:  $k = 1$ 
8: while  $h_{k,k-1} \neq 0$  do
9:    $u = Qv_k$ 
10:  Compute  $v_{k+1} = P_G u$  // Compute next Krylov vector
11:  for  $i = 1, \dots, k$  do // Modified Gram-Schmidt
12:     $h_{i,k} = \langle v_i, u \rangle$ 
13:     $v_{k+1} = v_{k+1} - h_{i,k} v_i$ 
14:   $h_{k+1,k} = \sqrt{\langle v_{k+1}, u \rangle}$ 
15:  if  $h_{k+1,k} \neq 0$  then
16:     $v_{k+1} = v_{k+1} / h_{k+1,k}$ 
17:   $x_{k+1} = V_{k+1}^\circ \Theta_{k+1}(H_{k+1,k}, h_{k+1,k})$  // Update solution estimate
18:   $k = k + 1$ 

```

---

streamlined. In particular, the estimate of the solution may often be expressed more succinctly using quantities computed elsewhere in the algorithm. We give an example of this in section 3.

Observe that Algorithm 2.2 (and those that will follow) only aim to find  $x_*$  satisfying (2.1) and not  $y_*$ . One way to obtain the complete solution is to compute an estimate of the  $y_*$  component of the solution only once a good approximation to  $x_*$  has been found. To do so, suppose that  $x$  has been obtained as an estimate of  $x_*$ . We may then compute

$$(2.15) \quad y_* = \arg \min_{y \in \mathbb{R}^m} \|A^T y - (c - Qx)\|_N$$

for some appropriate norm  $\|\cdot\|_N$ . If  $G$  were positive definite, the norm  $\|\cdot\|_{G^{-1}}$  would be suitable, and  $y_*$  may be found by solving

$$\begin{bmatrix} G & A^T \\ A & \end{bmatrix} \begin{bmatrix} w_* \\ y_* \end{bmatrix} = \begin{bmatrix} c - Qx \\ \mathbf{0} \end{bmatrix}$$

involving the matrix  $K_G$  used in (2.13). The same is true using the dual to the seminorm  $\sqrt{\langle s, Gs \rangle}$  defined on the manifold  $As = \mathbf{0}$  under the more general Assumption 2.2 on  $G$  (see Conn et al., 2000, section 2.2; and Gould, Hribar, and Nocedal, 2001, section 6).

Principle 2.2 may equally be applied to the Lanczos biorthogonalization process or its transpose-free variant to obtain corresponding full-space statements.

**2.4. Constraint-preconditioned variants.** In section 2.2 we showed how to form a  $G$ -orthogonal basis for the Krylov space generated from  $Q$  and  $c$ , using the Arnoldi process (Algorithm 2.1). This then lead us to an iteration (Algorithm 2.2) involving  $Q$  and the constraint preconditioner  $K_G$  in the full space. One might naturally ask whether it is possible to bypass this two-stage approach, and instead to



apply the preconditioned processes of section 2.2 directly to (2.1), with a constraint preconditioner of the form (2.13). Note that we temporarily ignore the fact that such a “preconditioner” is indefinite (and thus formally not applicable) and thus flout conventional wisdom; we simply write what the preconditioned processes would be if it were to be employed.

The following result establishes the equivalence between the projected methods derived above and standard preconditioned Krylov methods applied with a constraint preconditioner. Such a result has been known for some time in the case where  $\mathbf{Q}$  is symmetric and positive definite (see, e.g., Lukšan and Vlček, 1998; Perugia, Simoncini, and Arioli, 1999; Rozložník and Simoncini, 2002). Here we generalize the standard proofs to allow for general matrices.

**THEOREM 2.2.** *Algorithm 2.1 applied to (2.1) with an initial guess of the form  $[\mathbf{x}_0^\top, \mathbf{y}_0^\top]^\top$  satisfying  $\mathbf{A}\mathbf{x}_0 = \mathbf{0}$  and using the constraint preconditioner (2.13) generates Krylov vectors  $\mathbf{v}_k = [v_{k,x}^\top, v_{k,y}^\top]^\top$  such that at each iteration  $k$ ,  $\mathbf{v}_{k,x}$  is the Krylov vector  $\mathbf{v}_k$  generated by Algorithm 2.2 with initial guess  $\mathbf{x}_0$  at iteration  $k$ . The temporary vector  $\mathbf{u}$  has the form  $[(\mathbf{u} + \mathbf{A}^\top \mathbf{w})^\top, \mathbf{0}^\top]^\top$  for some vector  $\mathbf{w}$ , where  $\mathbf{u}$  is the corresponding temporary vector generated by Algorithm 2.2. In addition, the Hessenberg matrix generated is identical to that generated by Algorithm 2.2.*

*Proof.* Suppose Algorithm 2.1 is initialized with a starting guess of the form  $[\mathbf{x}_0^\top, \mathbf{y}_0^\top]^\top$  for which  $\mathbf{x}_0 \in \text{Null}(\mathbf{A})$ . Line 1 of Algorithm 2.1 reads

$$\mathbf{u} = \begin{bmatrix} \mathbf{c} \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{Q} & \mathbf{A}^\top \\ \mathbf{A} & \end{bmatrix} \begin{bmatrix} \mathbf{x}_0 \\ \mathbf{y}_0 \end{bmatrix} = \begin{bmatrix} \mathbf{c} - \mathbf{Q}\mathbf{x}_0 + \mathbf{A}^\top \mathbf{y}_0 \\ \mathbf{0} \end{bmatrix}$$

because  $\mathbf{A}\mathbf{x}_0 = \mathbf{0}$  so that the initial Krylov vector is given as the solution of

$$\begin{bmatrix} \mathbf{G} & \mathbf{A}^\top \\ \mathbf{A} & \end{bmatrix} \begin{bmatrix} v_{1,x} \\ v_{1,y} \end{bmatrix} = \begin{bmatrix} \mathbf{c} + \mathbf{Q}\mathbf{x}_0 + \mathbf{A}^\top \mathbf{y}_0 \\ \mathbf{0} \end{bmatrix}$$

and this shows that  $v_{1,x}$  is identical to the vector  $\mathbf{v}_1$  given at line 2 of Algorithm 2.2 since the term  $\mathbf{A}^\top \mathbf{y}_0$  does not impact the component  $v_{1,x}$  of the solution—see the second property of Theorem 2.1. In particular,  $v_{1,x}$  lies in the nullspace of  $\mathbf{A}$ . At line 3 of Algorithm 2.1, we compute

$$h_{1,0} = \sqrt{\langle v_{1,x}, \mathbf{c} - \mathbf{Q}\mathbf{x}_0 + \mathbf{A}^\top \mathbf{y}_0 \rangle + \langle v_{1,y}, \mathbf{0} \rangle} = \sqrt{\langle \mathbf{v}_1, \mathbf{c} - \mathbf{Q}\mathbf{x}_0 \rangle},$$

which is identical to  $h_{1,0}$  computed at line 3 of Algorithm 2.2.

Assuming now that  $\mathbf{v}_{k,x}$  lies in the nullspace of  $\mathbf{A}$  and coincides with  $\mathbf{v}_k$ , we establish by recursion that  $\mathbf{v}_{k+1,x}$  also lies in the nullspace of  $\mathbf{A}$  and coincides with  $\mathbf{v}_{k+1}$ .

The vector  $\mathbf{u}$  computed at line 8 of Algorithm 2.1 is

$$\mathbf{u} = \begin{bmatrix} \mathbf{Q} & \mathbf{A}^\top \\ \mathbf{A} & \end{bmatrix} \begin{bmatrix} v_{k,x} \\ v_{k,y} \end{bmatrix} = \begin{bmatrix} \mathbf{Q}\mathbf{v}_{k,x} + \mathbf{A}^\top v_{k,y} \\ \mathbf{0} \end{bmatrix}$$

and the first component of  $\mathbf{u}$  is  $\mathbf{u} + \mathbf{A}^\top v_{k,y}$ , where  $\mathbf{u}$  is computed at line 9 of Algorithm 2.2. Next,  $\mathbf{v}_{k+1}$  is given as the solution of

$$\begin{bmatrix} \mathbf{G} & \mathbf{A}^\top \\ \mathbf{A} & \end{bmatrix} \begin{bmatrix} v_{k+1,x} \\ v_{k+1,y} \end{bmatrix} = \begin{bmatrix} \mathbf{u} + \mathbf{A}^\top v_{k,y} \\ \mathbf{0} \end{bmatrix}.$$

This shows that  $\mathbf{v}_{k+1,x} = \mathbf{v}_{k+1}$ . Moreover,  $\mathbf{A}\mathbf{v}_{k+1,x} = \mathbf{0}$ .

Finally, Algorithm 2.1 computes

$$h_{i,k} = \langle v_{i,x}, \mathbf{Q}v_{k,x} + \mathbf{A}^\top v_{k,y} \rangle + \langle v_{i,y}, \mathbf{0} \rangle = \langle v_{i,x}, \mathbf{Q}v_{k,x} \rangle = \langle \mathbf{v}_i, \mathbf{u} \rangle = \mathbf{h}_{i,k},$$

where we used the fact that  $\mathbf{A}v_{i,x} = 0$  for all  $i \leq k$ . Similarly, we can show that  $h_{k+1,k} = \mathbf{h}_{k+1,k}$ .  $\square$

In exact arithmetic, the advantage of Algorithm 2.2 over Algorithm 2.1 with a constraint preconditioner is in terms of memory requirements. The vectors  $\mathbf{u}$  and  $\mathbf{v}_k$  generated have size  $n$  while the vectors  $u$  and  $v_k$  of Algorithm 2.1 have size  $n + m$ . The advantage of Algorithm 2.1 with constraint preconditioner is that the subvectors  $v_{k,y}$  may be used to recur approximations to the  $\mathbf{y}$  segment of the solution of (2.1). Depending on how the application of  $\mathbf{P}_G$  is implemented, however, the same approximations may be obtained and recurred in Algorithm 2.2.

The above equivalence applies to all Krylov methods deriving directly from the Arnoldi process, including GMRES and FOM, but also to methods deriving directly from the symmetric Lanczos process, including the conjugate gradient method, MINRES, and SYMMLQ (see Gould, Orban, and Rees, 2013). This illustrates the point that in the special context of the constraint-preconditioned symmetric Lanczos process and with a suitable initial guess, MINRES with the indefinite preconditioner (2.13) is a well-defined Krylov method.

*Example 2.1.* Projected GMRES Standard GMRES performs  $m$  iterations of Lanczos orthogonalization to obtain  $\mathbf{V}_{m+1}$  and  $\mathbf{H}_{m+1,m}$ , and then computes an approximation  $\mathbf{x}_{m+1}$  in the  $(m + 1)$ st Krylov subspace  $\mathbf{x}_{m+1} := \mathbf{V}_{m+1}\mathbf{z}_{m+1}$  by solving the linear least-squares problem

$$(2.16) \quad \underset{\mathbf{z}_{m+1}}{\text{minimize}} \quad \|\mathbf{h}_{1,0}e_1 - \mathbf{H}_{m+1,m}\mathbf{z}_{m+1}\|_2.$$

The preconditioned and projected GMRES algorithms are based on Algorithms 2.1 and 2.2, respectively. Besides Algorithm 2.2, the details of the projected GMRES algorithm are standard. The constraint-preconditioned GMRES algorithm defines its  $(m + 1)$ st approximation as

$$\begin{bmatrix} x_{m+1} \\ y_{m+1} \end{bmatrix} = \mathbf{V}_{m+1}\mathbf{z}_{m+1} = \begin{bmatrix} V_{x,m+1}\mathbf{z}_{m+1} \\ V_{y,m+1}\mathbf{z}_{m+1} \end{bmatrix}.$$

From Theorem 2.2 and (2.16), it is clear that the  $\mathbf{x}_{m+1}$  defined by the projected variant coincides with the  $x_{m+1}$  defined by the constraint-preconditioned variant.

Results similar in spirit to Theorem 2.2 may be stated to establish a formal equivalence between projected and constraint-preconditioned methods based on the Lanczos biorthogonalization process or its transpose-free variant (see Gould, Orban, and Rees, 2013). In particular, for any Krylov method deriving from the Lanczos biorthogonalization process, the  $x$  component of iterates generated by an appropriately initialized constraint-preconditioned variant of this method will coincide with the iterates generated by the projected variant of this method. Among others, this conclusion applies to the two-sided Lanczos method for linear systems, the biconjugate gradient algorithm (Bi-CG), and the quasi-minimum residual method (QMR). Just as for the Arnoldi process, such results have been known for some time in the case where  $\mathbf{Q}$  is symmetric and positive definite (see, e.g., Lukšan and Vlček, 1998; Perugia, Simoncini, and Arioli, 1999; Rozložník and Simoncini, 2002).

**3. Example. Projected MINRES.** In this section, we consider the specific example of the projected MINRES algorithm, in which approximate solutions  $x_k = V_k z_k$  are chosen so as to minimize the norm of the residual  $r_k := c - Qx_k$  in the norm defined by  $G^{-1}$ . Recall that  $c = Z^T c$ ,  $Q = Z^T Q Z$ , and  $G = Z^T G Z$  is positive definite. The quantity minimized by the projected MINRES at each iteration may thus be written

$$(3.1) \quad \|r_k\|_{G^{-1}}^2 = \langle r_k, (Z^T G Z)^{-1} r_k \rangle.$$

We now establish directly that standard MINRES applied with a constraint preconditioner also minimizes (3.1), and therefore that MINRES with the indefinite preconditioner (2.13) is a well-defined Krylov method.

Theorem 2.2 guarantees that with appropriate initial conditions, the approximate solution  $(x_k, y_k)$  generated at the  $k$ th iteration of the constraint-preconditioned MINRES is such that  $x_k \in \text{Null}(A)$ . Consider the residual

$$r_k := \begin{bmatrix} c \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} Q & A^T \\ A & \end{bmatrix} \begin{bmatrix} x_k \\ y_k \end{bmatrix} = \begin{bmatrix} c - Qx_k - A^T y_k \\ \mathbf{0} \end{bmatrix}.$$

It is tempting to claim that it is this residual that is minimized in the norm defined by the inverse of the preconditioner (2.13). Unfortunately, the preconditioner is indefinite. The appropriate interpretation of this claim is to consider the *seminorm* associated with the preconditioner and defined by

$$(3.2) \quad \|r_k\|_{[G]}^2 := \langle r_k, s_k \rangle, \quad \text{where} \quad \begin{bmatrix} G & A^T \\ A & \end{bmatrix} \begin{bmatrix} s_k \\ t_k \end{bmatrix} = \begin{bmatrix} r_k \\ \mathbf{0} \end{bmatrix},$$

i.e.,

$$(3.3) \quad \|r_k\|_{[G]}^2 = \langle c - Qx_k - A^T y_k, s_k \rangle = \langle c - Qx_k, s_k \rangle.$$

Note that  $\|r_k\|_{[G]}$  measures deviation of  $r_k$  from the range space of  $A^T$  and vanishes if and only if  $r_k$  is orthogonal to the nullspace of  $A$ . In effect,  $\|\cdot\|_{[G]}$  defines a norm on the nullspace of  $A$ . Such seminorms have been used in optimization contexts (Conn et al., 2000). Observe from (3.2) that  $s_k \in \text{Null}(A)$  and therefore,  $s_k = Z s_k$  for some  $s_k$ . The first block equation of (3.2) premultiplied by  $Z^T$  then yields

$$s_k = (Z^T G Z)^{-1} Z^T (c - Qx_k) = (Z^T G Z)^{-1} r_k.$$

Introducing this expression of  $s_k$  into (3.3), we finally obtain

$$\|r_k\|_{[G]}^2 = \langle Z^T (c - Qx_k), s_k \rangle = \langle r_k, (Z^T G Z)^{-1} r_k \rangle = \|r_k\|_{G^{-1}}^2,$$

which coincides with (3.1). On substituting  $r_k = Z^T r_k$  in this last identity, we also see that

$$\|r_k\|_{[G]}^2 = \langle r_k, Z(Z^T G Z)^{-1} Z^T r_k \rangle = \langle r_k, P_G(r_k) \rangle = \|P_G(r_k)\|_2^2.$$

We conclude that the projected MINRES algorithm minimizes the Euclidean norm of the projected residual, i.e., its seminorm induced by the constraint preconditioner. Since the spaces over which this quantity are minimized are the same in both methods (Rozložník and Simoncini, 2002), projected MINRES and standard MINRES applied with a constraint preconditioner are equivalent in exact arithmetic.

Similar conclusions can be drawn about the projected SYMMLQ algorithm, which is also well defined in the presence of a constraint preconditioner.

As with the projected conjugate gradient method, the numerical stability of both projected MINRES and standard MINRES applied with a constraint preconditioner is dependent on keeping the components  $\mathbf{x}_k$  in the nullspace of  $\mathbf{A}$ . While this is always true in exact arithmetic, the accumulation of rounding errors can quickly give  $\mathbf{x}_k$  a nonnegligible component in the range of  $\mathbf{A}^\top$  which may cause the method to break down.

To help to ameliorate this effect Gould, Hribar, and Nocedal (2001) suggest that the constraint preconditioner be applied exactly, which is achieved by not only solving the preconditioned system with a direct method, but by applying one (or more) steps of iterative refinement to the approximate solution obtained. This was suggested in the context of projected CG, but the technique is still worthwhile here.

*Example 3.1.* We take the matrix and preconditioner formed in MATLAB by the commands.

---

```

1  n = 100; m = 75;
2  Q = rand(n,n);
3  Q = Q + Q' + 5*eye(n);
4  A = rand(m,n);
5  K = [Q A'; A zeros(m,m)];
6  G = diag(abs(diag(Q)));

```

---

That is, we consider a random saddle-point matrix with a diagonally dominant leading block and take the constraint preconditioner, where  $\mathbf{G}$  is the positive diagonal matrix whose nonzero entries coincide with those of  $\mathbf{Q}$ .

Example 3.1 is chosen so that the approximate leading block  $\mathbf{G}$  is a good approximation of the actual  $\mathbf{Q}$ . In the next example this is not the case.

*Example 3.2.* Here we modify the matrix from Example 3.1 by setting  $\mathbf{Q} = \mathbf{Q} - 5 * \mathbf{eye}(5)$ , and taking the equivalent (1,1) block in the constraint preconditioner.

We apply projected MINRES (denoted PPMINRES) and standard MINRES with a constraint preconditioner to the linear systems described in Examples 3.1 and 3.2. The algorithms are applied both with and without a step of iterative refinement. The results are reported as convergence curves in Figure 1.

From Figure 1, we see that—as was the case with projected CG (Gould, Hribar, and Nocedal, 2001)—it is advisable to apply these methods with iterative refinement.

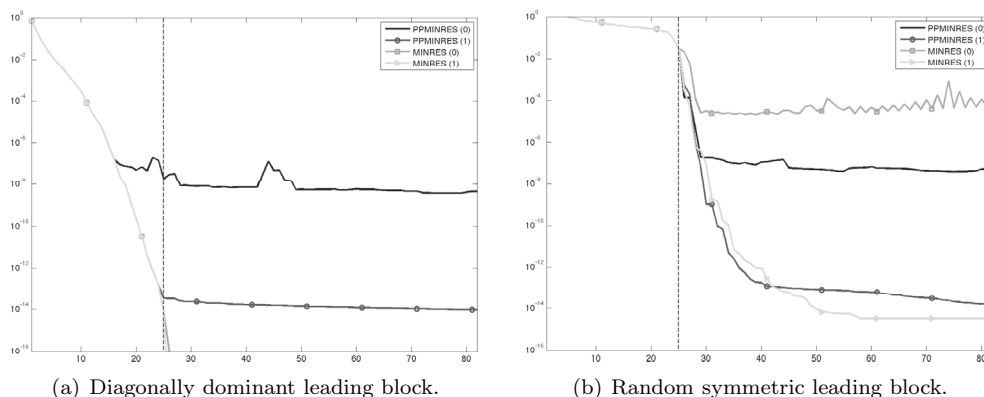


FIG. 1. Convergence curves for MINRES and PPMINRES. The number in brackets denotes the number of steps of iterative refinement used. The vertical dashed line shows where the algorithm would converge in exact arithmetic.

Our experiments suggest that a single step of iterative refinement is sufficient. If the  $(1, 1)$  block of the preconditioner is a good approximation to  $\mathbf{Q}$ , then it seems that standard MINRES alone does very well. However, when the approximation is not so good, as in Example 3.2, this becomes the worst-performing method in terms of maximum attainable accuracy.

We note that Fong and Saunders (2012) argue that it is often advantageous to use MINRES over CG even for symmetric positive definite systems of equations. Since we showed that MINRES applied with a constraint preconditioner is a well-defined method, there may be advantages to using projected MINRES in the place of projected CG even for cases where  $\mathbf{Q}$  is symmetric positive definite.

**4. Discussion.** We have considered the families of processes that form the basis of the great majority of Krylov methods. There are other types of processes that also possess projected variants although they do not strictly qualify as Krylov methods. An example is the tridiagonalization process described by Saunders, Simon, and Yip (1988) on which the iterative methods USYMLQ and USYMQR are based. Initialized with  $v_0 = w_0 = 0$  and arbitrary unit vectors  $v_1$  and  $w_1$ , this process generates sequences  $\{v_k\}$  and  $\{w_k\}$  according to

$$(4.1a) \quad t_{k+1,k}v_{k+1} = Qw_k - t_{k,k}v_k - t_{k-1,k}v_{k-1},$$

$$(4.1b) \quad t_{k,k+1}w_{k+1} = Q^T v_k - t_{k,k}w_k - t_{k,k-1}w_{k-1},$$

where  $t_{k,k} := \langle v_k, Qw_k \rangle$  and the off-diagonal elements  $t_{k+1,k}$  and  $t_{k,k+1}$  are chosen to normalize  $v_{k+1}$  and  $w_{k+1}$ . On rearranging the above, the process is characterized by the identities

$$QW_k = V_k T_k + t_{k+1,k}v_{k+1}e_k^T,$$

$$Q^T V_k = W_k T_k + t_{k,k+1}w_{k+1}e_k^T,$$

where  $T_k$  is a  $k$ -by- $k$  tridiagonal matrix with positive off-diagonal elements. It is possible to show that both sequences  $\{v_k\}$  and  $\{w_k\}$  are orthonormal and they are mutually conjugate in the sense that  $\langle v_j, Qw_k \rangle = \langle w_j, Qv_k \rangle = 0$  for  $k < j - 1$  (Saunders, Simon, and Yip, 1988, Theorem 1). The above relations differ considerably from those arising from biorthogonalization processes but are reminiscent of the Golub and Kahan (1965) bidiagonalization process. Though the above process does not generate Krylov spaces but somewhat larger spaces, it is possible to derive a projected variant applicable directly to (2.1) in the same way as in section 2.3.

Orban (2014) applies Principles 2.1 and 2.2 to the Golub and Kahan (1965) process to derive projected variants of iterative methods for linear least-squares problems such as LSQR (Paige and Saunders, 1982) and LSMR (Fong and Saunders, 2011). The resulting methods are suited to linear least-squares problems with linear equality constraints and are intimately related to full-space methods such as MINRES applied to the symmetric system representing the optimality conditions of such problems.

Certain saddle-point systems, such as those arising from the discretization of stabilized Navier–Stokes flow problems, have the form

$$\begin{bmatrix} \mathbf{Q} & \mathbf{A}^T \\ \mathbf{A} & -\mathbf{C} \end{bmatrix} \begin{bmatrix} \mathbf{x}_* \\ \mathbf{y}_* \end{bmatrix} = \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix},$$

where  $\mathbf{C}$  is typically symmetric and positive semidefinite (Elman, Silvester, and Wathen, 2005). Assuming more generally that  $\mathbf{C}$  may be decomposed as  $\mathbf{EDE}^T$ , where

$D$  is nonsingular, and introducing  $\mathbf{w} := -D\mathbf{E}^\top \mathbf{y}$  as suggested by Dollar et al. (2006), such systems may be equivalently reformulated as

$$\begin{bmatrix} Q & & A^\top \\ & D^{-1} & E^\top \\ A & E & \end{bmatrix} \begin{bmatrix} \mathbf{x}_* \\ \mathbf{w}_* \\ \mathbf{y}_* \end{bmatrix} = \begin{bmatrix} \mathbf{a} \\ \mathbf{0} \\ \mathbf{b} \end{bmatrix}.$$

The latter system has the form (1.1) and the methods proposed in this paper are applicable.

In the case where  $Q$  is symmetric positive definite on the nullspace of  $A$ ,  $\langle \cdot, \cdot \rangle_{K_Q}$  defines a semidefinite bilinear form on the space  $\mathcal{L} := \{[\mathbf{x}; \mathbf{y}] \text{ s.t. } \mathbf{x} \in \ker(A)\}$ , and hence induces a seminorm on  $\mathcal{L}$ . If we restrict  $\mathcal{L}$  to the subset where  $\mathbf{y} = \mathbf{0}$ , then the bilinear form above becomes an inner product. The matrix

$$\mathcal{A} := \begin{bmatrix} G & A^\top \\ A & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} Q & A^\top \\ A & \mathbf{0} \end{bmatrix}$$

is trivially self-adjoint in this bilinear form, and since the space  $\mathcal{L}$  is invariant under multiplication with  $\mathcal{A}$ , projected conjugate gradients can be thought of as standard conjugate gradients, with a suitable starting vector in  $\mathcal{L}$ , in this nonstandard (indefinite) inner product. The extensions to other Krylov methods, as outlined in the preceding sections, can be interpreted similarly where appropriate.

Finally, we note that it appears as if  $A$  must be available as an explicit matrix to factorize (2.13). There are other ways to compute projections. One of them is to solve a linear least-squares problem of the form (2.15) using an appropriate metric, which only requires  $A$  to be available as an operator. It remains important, however, that such projections be computed accurately and this can mean that an iterative solver must be supplied with stringent stopping conditions.

**5. Conclusion.** We have demonstrated systematic principles to show how any Krylov subspace method may be applied as a projected method, which acts purely in the nullspace. In particular—contrary to popular belief—methods for symmetric but indefinite matrices, such as MINRES, may be applied as projected methods.

Furthermore, we have shown that any Krylov subspace method applied with a constraint preconditioner is numerically equivalent to a projected Krylov method. A comparison of the relative merits of the different methods when applied with a constraint preconditioner would be interesting to perform, but is beyond the scope of the current work. Although the projected and standard methods are equivalent in exact arithmetic they have been shown to behave differently in the presence of rounding errors; a detailed investigation of this phenomenon would be another interesting avenue for further work.

#### REFERENCES

- W. E. ARNOLDI (1951), *The principle of minimized iterations in the solution of the matrix eigenvalue problem*, Quart. Appl. Math., 9, pp. 17–29.
- M. BENZI, G. H. GOLUB, AND J. LIESEN (2005), *Numerical solution of saddle point problems*, Acta Numer., 14, pp. 1–137.
- C. BREZINSKI AND M. REDIVO-ZAGLIA (1998), *Transpose-free Lanczos-type algorithms for nonsymmetric linear systems*, Numer. Algorithms, 17, pp. 67–103.
- T. F. CHAN, L. DE PILLIS, AND H. VAN DER VORST (1998), *Transpose-free formulations of Lanczos-type methods for nonsymmetric linear systems*, Numer. Algorithms, 17, pp. 51–66.

- T. F. COLEMAN (1994), *Linearly constrained optimization and projected preconditioned conjugate gradients*, in Proceedings of the Fifth SIAM Conference on Applied Linear Algebra, J. G. Lewis, ed., SIAM, Philadelphia, pp. 118–122.
- A. R. CONN, N. I. M. GOULD, D. ORBAN, AND P. TOINT (2000), *A primal-dual trust-region algorithm for non-convex nonlinear programming*, Math. Program. B, 87, pp. 215–249.
- H. S. DOLLAR AND A. J. WATHEN (2006), *Approximate factorization constraint preconditioners for saddle-point matrices*, SIAM J. Sci. Comput., 27, pp. 1555–1572.
- H. S. DOLLAR, N. I. M. GOULD, W. H. A. SCHILDERS, AND A. J. WATHEN (2006), *Implicit-factorization preconditioning and iterative solvers for regularized saddle-point systems*, SIAM J. Matrix Anal. Appl., 28, pp. 170–189.
- H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN (2005), *Finite-Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, Oxford University Press, Oxford.
- D. C.-L. FONG AND M. SAUNDERS (2011), *LSMR: An iterative algorithm for sparse least-squares problems*, SIAM J. Sci. Comput., 33, pp. 2950–2971.
- D. C.-L. FONG AND M. A. SAUNDERS, *CG versus MINRES: An empirical comparison*, SQU J. Sci., 17 (2012), pp. 44–62.
- R. FREUND, M. GUTKNECHT, AND N. NACHTIGAL (1993), *An implementation of the look-ahead Lanczos algorithm for non-Hermitian matrices*, SIAM J. Sci. Comput., 14, pp. 137–158.
- G. GOLUB AND W. KAHAN (1965), *Calculating the singular values and pseudo-inverse of a matrix*, J. Soc. Indust. Appl. Math. Ser. B Numer. Anal., 2, pp. 205–224.
- N. I. M. GOULD (1985), *On practical conditions for the existence and uniqueness of solutions to the general equality quadratic-programming problem*, Math. Program., 32, pp. 90–99.
- N. I. M. GOULD, M. E. HRIBAR, AND J. NOCEDAL (2001), *On the solution of equality constrained quadratic problems arising in optimization*, SIAM J. Sci. Comput., 23, pp. 1376–1395.
- N. I. M. GOULD, D. ORBAN, AND T. REES (2013), *Projected Krylov Methods for Saddle-Point Systems*, Technical report RAL-P-2013-006, Rutherford Appleton Laboratory, Chilton, Oxfordshire, England.
- A. GREENBAUM (1997), *Iterative Methods for Solving Linear Systems*, Frontiers Appl. Math. 17, SIAM, Philadelphia.
- C. KELLER, N. I. M. GOULD, AND A. J. WATHEN (2000), *Constraint preconditioning for indefinite linear systems*, SIAM J. Matrix Anal. Appl., 21, pp. 1300–1317.
- C. LANCZOS (1950), *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, J. Res. Nat. Bur. Stand. B, 45, pp. 225–280.
- L. LUKŠAN AND J. VLČEK (1998), *Indefinitely preconditioned inexact Newton method for large sparse equality constrained nonlinear programming problems*, Numer. Linear Algebra Appl., 5, pp. 219–247.
- D. ORBAN (2008), *Projected Krylov Methods for Unsymmetric Augmented Systems*, GERAD Technical report G-2008-46, Mathematics and Industrial Engineering Department, Ecole Polytechnique de Montreal, Montreal, Canada.
- D. ORBAN (2014), *The Projected Golub-Kahan Process for Constrained Linear Least-Squares Problems*, Technical report Cahier du GERAD G-2014-15, GERAD, Montréal, Canada.
- C. C. PAIGE AND M. A. SAUNDERS (1975), *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12, pp. 617–629.
- C. C. PAIGE AND M. A. SAUNDERS (1982), *LSQR: An algorithm for sparse linear equations and sparse least squares*, ACM Trans. Math. Software, 8, pp. 43–71.
- I. PERUGIA AND V. SIMONCINI (2000), *Block-diagonal and indefinite symmetric preconditioners for mixed finite element formulations*, Numer. Linear Algebra Appl., 7, pp. 585–616.
- I. PERUGIA, V. SIMONCINI, AND M. ARIOLI (1999), *Linear algebra methods in a mixed approximation of magnetostatic problems*, SIAM J. Sci. Comput., 21, pp. 1085–1101.
- B. T. POLYAK (1969), *The conjugate gradient method in extremal problems*, U.S.S.R. Comput. Math. Math. Phys., 9, pp. 94–112.
- M. ROZLOZNÍK AND V. SIMONCINI (2002), *Krylov subspace methods for saddle point problems with indefinite preconditioning*, SIAM J. Matrix Anal. Appl., 24, pp. 368–391.
- Y. SAAD (2003), *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia.
- M. A. SAUNDERS, H. D. SIMON, AND E. L. YIP (1988), *Two conjugate-gradient-type methods for unsymmetric linear equations*, SIAM J. Numer. Anal., 25, pp. 927–940.
- H. A. VAN DER VORST (2003), *Iterative Krylov Methods for Large Linear Systems*, Cambridge University Press, Cambridge.