

## ITERATIVE SOLUTION OF AUGMENTED SYSTEMS ARISING IN INTERIOR METHODS\*

ANDERS FORSGREN<sup>†</sup>, PHILIP E. GILL<sup>‡</sup>, AND JOSHUA D. GRIFFIN<sup>§</sup>

**Abstract.** Iterative methods are proposed for certain augmented systems of linear equations that arise in interior methods for general nonlinear optimization. Interior methods define a sequence of *KKT equations* that represent the symmetrized (but indefinite) equations associated with Newton’s method for a point satisfying the perturbed optimality conditions. These equations involve both the primal and dual variables and become increasingly ill-conditioned as the optimization proceeds. In this context, an iterative linear solver must not only handle the ill-conditioning but also detect the occurrence of KKT matrices with the wrong matrix inertia. A one-parameter family of equivalent linear equations is formulated that includes the KKT system as a special case. The discussion focuses on a particular system from this family, known as the “doubly augmented system,” that is positive definite with respect to both the primal and dual variables. This property means that a standard preconditioned conjugate-gradient method involving both primal and dual variables will either terminate successfully or detect if the KKT matrix has the wrong inertia. Constraint preconditioning is a well-known technique for preconditioning the conjugate-gradient method on augmented systems. A family of constraint preconditioners is proposed that provably eliminates the inherent ill-conditioning in the augmented system. A considerable benefit of combining constraint preconditioning with the doubly augmented system is that the preconditioner need not be applied exactly. Two particular “active-set” constraint preconditioners are formulated that involve only a subset of the rows of the augmented system and thereby may be applied with considerably less work. Finally, some numerical experiments illustrate the numerical performance of the proposed preconditioners and highlight some theoretical properties of the preconditioned matrices.

**Key words.** large-scale nonlinear programming, nonconvex optimization, interior methods, augmented systems, KKT systems, iterative methods, conjugate-gradient method, constraint preconditioning

**AMS subject classifications.** 49J20, 49J15, 49M37, 49D37, 65F05, 65K05, 90C30

**DOI.** 10.1137/060650210

**1. Introduction.** This paper concerns the formulation and analysis of preconditioned iterative methods for the solution of augmented systems of the form

$$(1.1) \quad \begin{pmatrix} H & -A^T \\ A & G \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix},$$

with  $A$  an  $m \times n$  matrix,  $H$  symmetric, and  $G$  symmetric positive semidefinite. These equations arise in a wide variety of scientific and engineering applications, where they are known by a number of different names, including “augmented systems,” “saddle-point systems,” “KKT systems,” and “equilibrium systems.” (The bibliography of the survey by Benzi, Golub, and Liesen [3] contains 513 related articles.) The main focus

---

\*Received by the editors January 17, 2006; accepted for publication (in revised form) March 5, 2007; published electronically August 22, 2007. The research of the second and third authors was supported by National Science Foundation grants DMS-9973276, CCF-0082100, and DMS-0511766.

<http://www.siam.org/journals/siopt/18-2/65021.html>

<sup>†</sup>Optimization and Systems Theory, Department of Mathematics, Royal Institute of Technology, SE-100 44 Stockholm, Sweden (andersf@kth.se). The research of this author was supported by the Swedish Research Council (VR).

<sup>‡</sup>Department of Mathematics, University of California, San Diego, La Jolla, CA 92093-0112 (pgill@ucsd.edu).

<sup>§</sup>Sandia National Laboratories, Livermore, CA 94551-9217 (jgriffi@sandia.gov). Part of this work was carried out during the Spring of 2003 while this author was visiting KTH with financial support from the Göran Gustafsson Foundation.

of this paper will be on the solution of augmented systems arising in interior methods for general constrained optimization, in which case (1.1) is the system associated with Newton's method for finding values of the primal and dual variables that satisfy the perturbed KKT optimality conditions (see, e.g., Wright [49] and Forsgren, Gill, and Wright [15]). In this context  $H$  is the Hessian of the Lagrangian,  $A$  is the constraint Jacobian, and  $G$  is diagonal.

Many of the benefits associated with the methods discussed in this paper derive from formulating the interior method so that the diagonal  $G$  is *positive definite*. We begin by presenting results for  $G$  positive definite and consider the treatment of systems with positive semidefinite and singular  $G$  in section 5. Throughout, for the case where  $G$  is positive definite, we denote  $G$  by  $D$  and rewrite (1.1) as an equivalent *symmetric* system  $Bx = b$ , where

$$(1.2) \quad B = \begin{pmatrix} H & -A^T \\ -A & -D \end{pmatrix} \quad \text{and} \quad b = \begin{pmatrix} b_1 \\ -b_2 \end{pmatrix},$$

with  $D$  positive definite and diagonal. We will refer to this symmetric system as the *KKT system*. (It is possible to symmetrize (1.1) in a number of different ways. The format (1.2) will simplify the linear algebra in later sections.) When  $D$  is nonsingular, it is well known that the augmented system is equivalent to the two smaller systems

$$(1.3) \quad (H + A^T D^{-1} A)x_1 = b_1 + A^T D^{-1} b_2 \quad \text{and} \quad x_2 = D^{-1}(b_2 - Ax_1),$$

where the system for  $x_1$  is known as the *condensed system*. It is less well known that another equivalent system is the *doubly augmented system*

$$(1.4) \quad \begin{pmatrix} H + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} b_1 + 2A^T D^{-1} b_2 \\ b_2 \end{pmatrix},$$

which has been proposed for use with direct factorization methods by Forsgren and Gill [16]. In this paper we investigate the properties of preconditioned iterative methods applied to system (1.2) directly or to the equivalent systems (1.3) and (1.4).

If the underlying optimization problem is not convex, the matrix  $H$  may be indefinite. The KKT matrix  $B$  of (1.2) is said to have *correct inertia* if the matrix  $H + A^T D^{-1} A$  is positive definite. This definition is based on the properties of the underlying optimization problem. Broadly speaking, the KKT system has correct inertia if the problem is locally convex (for further details see, e.g., Forsgren and Gill [16], Forsgren [18], and Griffin [32]). If the KKT matrix has correct inertia, then systems (1.2)–(1.4) have a common unique solution (see section 2).

**1.1. Properties of the KKT system.** The main issues associated with using iterative methods to solve KKT systems are (i) *termination control*, (ii) *inertia control*, and (iii) *inherent ill-conditioning*. The first of these issues is common to other applications where the linear system represents a linearization of some underlying nonlinear system of equations. Issues (ii) and (iii), however, are unique to optimization and will be the principal topics of this paper.

In the context of interior methods, the KKT system (1.2) is solved as part of a two-level iterative scheme. At the outer level, nonlinear equations that define the first-order optimality conditions are parameterized by a small positive quantity  $\mu$ . The idea is that the solution of the parameterized equations should approach the solution of the optimization problem as  $\mu \rightarrow 0$ . At the inner level, equations (1.2) represent the symmetrized Newton equations associated with finding a zero of the

perturbed optimality conditions for a given value of  $\mu$ . Although systems (1.2)–(1.4) have identical solutions, an iterative method will generally produce a different sequence of iterates in each case (see section 3 for a discussion of the equivalence of iterative solvers in this context). An iterative method applied to the augmented system (1.2) or the doubly augmented system (1.4) treats  $x_1$  and  $x_2$  as independent variables, which is appropriate in the optimization context because  $x_1$  and  $x_2$  are associated with independent quantities in the perturbed optimality conditions (i.e., the primal and dual variables). In contrast, an iterative solver for the condensed system (1.3) will generate approximations to  $x_1$  only, with the variables  $x_2$  being defined as  $x_2 = D^{-1}(b_2 - Ax_1)$ . This becomes an important issue when an *approximate* solution is obtained by *truncating* the iterations of the linear solver. During the early outer iterations, it is usually inefficient to solve the KKT system accurately, and it is better to accept an inexact solution that gives a residual norm that is less than some factor of the norm of the right-hand side (see, e.g., Dembo, Eisenstat, and Steihaug [7]). For the condensed system, the residual for the second block of equations will be zero regardless of the accuracy of  $x_1$ , which implies that termination must be based on the accuracy of  $x_1$  alone. It is particularly important for the solver to place equal weight on  $x_1$  and  $x_2$  when system (1.2) is being solved in conjunction with a primal-dual trust-region method (see Gertz and Gill [20] and Griffin [32]). The conjugate-gradient version of this method exploits the property that the norms of the  $(x_1, x_2)$  iterates increase monotonically (see Steihaug [44]). This property does not hold for  $(x_1, x_2)$  iterates generated for the condensed system.

If the KKT matrix does not have the correct inertia, the solution of (1.2) is not useful, and the optimization continues with an alternative technique based on either implicitly or explicitly modifying the matrix  $H$  (see, e.g., Toint [45], Steihaug [44], Gould et al. [30], Hager [33], and Griffin [32]). It is therefore important that the iterative solver is able to detect if  $B$  does not have correct inertia.

As the perturbation parameter  $\mu$  is reduced, the KKT systems become *increasingly ill-conditioned*. The precise form of this ill-conditioning depends on the formulation of the interior method, but a common feature is that some diagonal elements of  $D$  are big and some are small. (It is almost always possible to formulate an interior method that requires the solution of an *unsymmetric* system that does not exhibit inevitable ill-conditioning as  $\mu \rightarrow 0$ . This unsymmetric system could be solved using an unsymmetric solver such as GMRES or QMR. Unfortunately, this approach is unsuitable for general KKT systems because an unsymmetric solver is unable to determine if the KKT matrix has correct inertia.) In section 3 we consider a preconditioned conjugate-gradient (PCG) method that provably removes the inherent ill-conditioning. In particular, we define a one-parameter family of preconditioners related to the class of so-called *constraint preconditioners* proposed by Keller, Gould, and Wathen [34]. Several authors have used constraint preconditioners in conjunction with the conjugate-gradient method to solve the indefinite KKT system (1.2) with  $b_2 = 0$  and  $D = 0$  (see, e.g., Lukšan and Vlček [36], Gould, Hribar, and Nocedal [29], Perugia and Simoncini [40], and Bergamaschi, Gondzio, and Zilli [4]). Recently, Dollar [12] and Dollar et al. [11] have proposed constraint preconditioners for system (1.2) with no explicit inertial or diagonal condition on  $D$ , but a full row-rank requirement on  $A$  and the assumption that  $b_2 = 0$ .

Methods that require  $b_2 = 0$  must perform an initial projection step that effectively shifts the right-hand side to zero. The constraint preconditioner then forces the  $x_1$  iterates to lie in the null space of  $A$ . A disadvantage with this approach is that the

constraint preconditioner must be applied exactly if subsequent iterates are to lie in the null space. This limits the ability to perform approximate solves with the preconditioner, as is often required when the matrix  $A$  has a PDE-like structure that also must be handled using an iterative solver (see, e.g., Saad [41], Notay [37], Simoncini and Szyld [43], and Elman et al. [14]). In section 3 we consider preconditioners that do not require the assumption that  $b_2 = 0$ , and hence do not require an accurate solve with the preconditioner.

**1.2. A PCG method for the KKT system.** The goal of this paper is to formulate iterative methods that not only provide termination control and inertia control, but also eliminate the inevitable ill-conditioning associated with interior methods. All these features are present in an algorithm based on applying a PCG method to the doubly augmented system (1.4). This system is positive definite if the KKT matrix has correct inertia, and gives equal weight to  $x_1$  and  $x_2$  for early terminations. As preconditioner we use the constraint preconditioner

$$(1.5) \quad P = \begin{pmatrix} M + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix},$$

where  $M$  is an approximation of  $H$  such that  $M + A^T D^{-1} A$  is positive definite. The equations  $Pv = r$  used to apply the preconditioner are solved by exploiting the equivalence of the systems

$$(1.6a) \quad \begin{pmatrix} M + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} r_1 \\ r_2 \end{pmatrix},$$

$$(1.6b) \quad \begin{pmatrix} M & -A^T \\ -A & -D \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} r_1 - 2A^T D^{-1} r_2 \\ -r_2 \end{pmatrix}, \quad \text{and}$$

$$(1.6c) \quad (M + A^T D^{-1} A)v_1 = r_1 - A^T D^{-1} r_2, \quad v_2 = D^{-1}(r_2 - Av_1)$$

(see section 3). This allows us to compute the solution of (1.6a) by solving either (1.6b) or (1.6c). (The particular choice will depend on the relative efficiency of the methods available to solve the condensed and augmented systems.)

We emphasize that the doubly augmented systems are never formed or factored explicitly. The matrix associated with the doubly augmented equations (1.4) is used only as an operator to define products of the form  $v = Bu$ . As mentioned above, the equations (1.6a) that apply the preconditioner are solved using either (1.6b) or (1.6c). An important property of the method is that these equations also may be solved using an iterative method. (It is safe to use the augmented or condensed system for the preconditioner equations  $Pv = r$  because the inertia of  $P$  is guaranteed by the choice of  $M$  (see section 3).)

In section 4 we formulate and analyze two variants of the preconditioner (1.5) that exploit the asymptotic behavior of the elements of  $D$ . The use of these so-called *active-set preconditioners* may require significantly less work when the underlying optimization problem has more constraints than variables. In section 5, we consider the case where  $G$  is positive semidefinite and singular. Finally, in section 6, we present some numerical examples illustrating the properties of the proposed preconditioners.

**1.3. Notation and assumptions.** Unless explicitly indicated otherwise,  $\|\cdot\|$  denotes the vector two-norm or its subordinate matrix norm. The inertia of a real symmetric matrix  $A$ , denoted by  $\text{In}(A)$ , is the integer triple  $(a_+, a_-, a_0)$  giving the number of positive, negative, and zero eigenvalues of  $A$ . The spectrum of a (possibly unsymmetric) matrix  $A$  is denoted by  $\text{eig}(A)$ . As the analysis concerns matrices

with only real eigenvalues,  $\text{eig}(A)$  is regarded as an ordered set, with the least (i.e., “leftmost”) eigenvalue, denoted by  $\text{eig}_{\min}(A)$ , appearing first. The quantity  $\sigma_k(A)$  denotes the  $k$ th largest singular value of  $A$ . Given a positive-definite  $A$ , the unique positive-definite  $X$  such that  $X^2 = A$  is denoted by  $A^{1/2}$ . Given vectors  $x_1$  and  $x_2$ , the column vector consisting of the elements of  $x_1$  augmented by the elements of  $x_2$  is denoted by  $(x_1, x_2)$ .

When  $\mu$  is a positive scalar such that  $\mu \rightarrow 0$ , the notation  $p = O(\mu)$  means that there exists a constant  $K$  such that  $|p| \leq K\mu$  for all  $\mu$  sufficiently small. For a positive  $p$ ,  $p = \Omega(1/\mu)$  implies that there exists a constant  $K$  such that  $1/p \leq K\mu$  for all  $\mu$  sufficiently small. In particular,  $p = O(1)$  means that  $|p|$  is bounded, and, for a positive  $p$ ,  $p = \Omega(1)$  means that  $p$  is bounded away from zero. For a positive  $p$ , the notation  $p = \Theta(1)$  is used for the case where both  $p = O(1)$  and  $p = \Omega(1)$ , so that  $p$  remains bounded and is bounded away from zero as  $\mu \rightarrow 0$ .

As discussed in section 1.1, we are concerned with solving a sequence of systems of the form (1.2), where the matrices  $A$ ,  $H$ , and  $D$  depend implicitly on  $\mu$ . In particular,  $A$  and  $H$  are first and second derivatives evaluated at a point depending on  $\mu$ , and  $D$  is an explicit function of  $\mu$ . The notation defined above allows us to characterize the properties of  $H$ ,  $A$ , and  $D$  in terms of their behavior as  $\mu \rightarrow 0$ . Throughout the analysis, it is assumed that the following properties hold:

(A<sub>1</sub>)  $\|H\|$  and  $\|A\|$  are both  $O(1)$ .

(A<sub>2</sub>) The row indices of  $A$  may be partitioned into disjoint subsets  $\mathcal{S}$ ,  $\mathcal{M}$ , and  $\mathcal{B}$  such that  $d_{ii} = O(\mu)$  for  $i \in \mathcal{S}$ ,  $d_{ii} = \Theta(1)$  for  $i \in \mathcal{M}$ , and  $d_{ii} = \Omega(1/\mu)$  for  $i \in \mathcal{B}$ .

(A<sub>3</sub>) If  $A_{\mathcal{S}}$  is the matrix of rows of  $A$  with indices in  $\mathcal{S}$  and  $r = \text{rank}(A_{\mathcal{S}})$ , then  $r$  remains constant as  $\mu \rightarrow 0$  and  $\sigma_r(A_{\mathcal{S}}) = \Theta(1)$ .

The second assumption reflects the fact that for  $\mu$  sufficiently small, some diagonal elements of  $D$  are “small,” some are “medium,” and some are “big.”

It is often the case in practice that the equations and variables corresponding to unit rows of  $A$  are eliminated directly from the KKT system. This elimination creates no additional nonzero elements and provides a smaller “partially condensed” system with an  $\Omega(1/\mu)$  diagonal term added to  $H$ . It will be shown that preconditioners for both the full and partially condensed KKT systems depend on the eigenvalues of the same matrix (see Lemmas 3.4 and 3.5). It follows that our analysis also applies to preconditioners defined for the partially condensed system.

**2. A parameterized system of linear equations.** In this section, it is shown how the indefinite KKT system (1.2) may be embedded in a family of equivalent linear systems, parameterized by a scalar  $\nu$ . This parameterization facilitates the simultaneous analysis of the three systems (1.2)–(1.4).

**DEFINITION 2.1** (the parameterized system). *Let  $\nu$  denote a scalar. Associated with the KKT equations  $Bx = b$  of (1.2), we define the parameterized equations  $B(\nu)x = b(\nu)$ , with*

$$B(\nu) = \begin{pmatrix} H + (1 + \nu)A^T D^{-1} A & \nu A^T \\ \nu A & \nu D \end{pmatrix} \quad \text{and} \quad b(\nu) = \begin{pmatrix} b_1 + (1 + \nu)A^T D^{-1} b_2 \\ \nu b_2 \end{pmatrix},$$

where  $H$  is symmetric and  $D$  is positive definite and diagonal.

The following proposition states the equivalence of the KKT system (1.2) and the parameterized system of Definition 2.1.

**PROPOSITION 2.2** (equivalence of the parameterized systems). *Let  $\nu$  denote a scalar parameter. If  $\nu \neq 0$ , then the system  $Bx = b$  of (1.2) and the system  $B(\nu)x =$*

$b(\nu)$  of Definition 2.1 are equivalent, i.e., (1.2) has a solution  $(x_1, x_2)$  if and only if  $(x_1, x_2)$  is a solution to  $B(\nu)x = b(\nu)$ . If  $\nu = 0$ , then (1.2) has a solution  $(x_1, x_2)$  if and only if  $x_1$  is a solution of  $B(0)x = b(0)$  and  $x_2 = D^{-1}(b_2 - Ax_1)$ .

We are particularly interested in the parameterized system  $B(\nu)x = b(\nu)$  for the values  $\nu = -1$ ,  $\nu = 0$ , and  $\nu = 1$ . If  $\nu = -1$ , we obtain the symmetric KKT system (1.2). If  $\nu = 0$ , we obtain the condensed system

$$(2.1) \quad (H + A^T D^{-1} A)x_1 = b_1 + A^T D^{-1} b_2 \quad \text{with } x_2 \text{ arbitrary.}$$

In this case,  $x_2$  does not appear in the augmented system and must be computed as  $x_2 = D^{-1}(b_2 - Ax_1)$ . Finally, if  $\nu = 1$ , we obtain the doubly augmented system

$$(2.2) \quad \begin{pmatrix} H + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} b_1 + 2A^T D^{-1} b_2 \\ b_2 \end{pmatrix}.$$

The next result follows from Lemma A.1 of the appendix and gives the inertia of  $B(\nu)$  as a function of  $\nu$ .

**PROPOSITION 2.3** (inertia of the parameterized system). *For the matrix  $B(\nu)$  of Definition 2.1, it holds that*

- (i)  $\text{In}(B(\nu)) = \text{In}(H + A^T D^{-1} A) + (m, 0, 0)$  if  $\nu > 0$ ;
- (ii)  $\text{In}(B(\nu)) = \text{In}(H + A^T D^{-1} A) + (0, m, 0)$  if  $\nu < 0$ ; and
- (iii)  $\text{In}(B(\nu)) = \text{In}(H + A^T D^{-1} A) + (0, 0, m)$  if  $\nu = 0$ .

This proposition implies that the inertia of  $B(\nu)$  may be determined from the inertia of  $H + A^T D^{-1} A$  and the sign of  $\nu$ . In particular, the result gives the inertia of the three alternate systems in the situation where the inertia of  $B$  is correct and the solver can be allowed to continue solving the system. If the inertia of  $B$  is correct, then (i) the doubly augmented system is positive definite; (ii) the KKT system has  $n$  positive eigenvalues and  $m$  negative eigenvalues; and (iii) the condensed system is positive definite (when regarded as a system involving only  $x_1$ ).

Proposition 2.3 implies that it is not worth applying a conjugate-gradient method to the general indefinite KKT system (1.2) because this method is unable to estimate the number of negative eigenvalues of an indefinite matrix. In contrast, the conjugate-gradient method is appropriate for both the doubly augmented system and the condensed system because indefiniteness is immediately indicated by the occurrence of a negative value of  $p_j^T C p_j$ , where  $p_i$  is a conjugate direction and  $C$  is either the doubly augmented matrix or the matrix for the condensed system. In other words, the occurrence of a negative value of  $p_j^T C p_j$  indicates that the inertia of the system is incorrect and the search for a solution of (1.2) should be abandoned.

**3. Constraint preconditioning for the linear equations.** The rate of convergence of the conjugate-gradient method may be accelerated by choosing an appropriate symmetric positive-definite preconditioner of the form  $P = R^T R$ , and applying the conjugate-gradient method to the preconditioned system  $R^{-T} B R^{-1} R x = R^{-T} b$ . As is well known, the computations may be arranged so that the preconditioner is applied by solving systems of the form  $P v = r$ . It is the eigenvalues of the preconditioned matrix  $R^{-T} B R^{-1}$  that determine the rate of convergence. As  $\text{eig}(R^{-T} B R^{-1}) = \text{eig}(R^{-1} R^{-T} B R^{-1} R) = \text{eig}(P^{-1} B)$ , the analysis may be written in terms of  $P^{-1} B$  without regard to  $R$ . However, it must be emphasized that  $P$  must be symmetric positive definite for the standard PCG method to be well defined.

Several authors have suggested constraint preconditioners for (1.2) and (1.3), in which  $H$  is replaced by a ‘‘simpler’’ approximation matrix  $M$  such that  $M + A^T D^{-1} A$  is

positive definite (see, e.g., Keller, Gould, and Wathen [34] and Bergamaschi, Gondzio, and Zilli [4]).

Under certain circumstances, the PCG method may be applied to all three systems and will give identical results in exact arithmetic.

PROPOSITION 3.1. *Assume that  $H + A^T D^{-1} A$  is positive definite. Consider the PCG method applied to the KKT system (1.2), the condensed system (1.3), and the doubly augmented system (1.4) with preconditioners*

$$(3.1) \quad \begin{pmatrix} M & -A^T \\ -A & -D \end{pmatrix}, \quad M + A^T D^{-1} A, \quad \text{and} \quad \begin{pmatrix} M + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix},$$

respectively. If  $b_2 = 0$ , then the PCG method generates the same sequence of iterates for all three systems (where an  $x_2$  iterate for the condensed system is defined as the product of  $-D^{-1} A$  and the  $x_1$  iterate).

The first preconditioner of (3.1) is not positive definite, which implies that it does not fit within the conventional PCG framework. However, Proposition 3.1 implies that the PCG method may be applied safely to the KKT system (1.2) in the special situation where  $b_2 = 0$ . The result is a *projected* PCG method that can be shown to be formally equivalent to the standard method applied to the condensed system; see, e.g., Lukšan and Vlček [36] and Gould, Hribar, and Nocedal [29].

The condition  $b_2 = 0$  may be achieved by choosing a special initial point  $y$ . In particular, consider the point  $(y_1, y_2)$  such that  $Ay_1 + Dy_2 = b_2$ , and the appropriate preconditioner is used for each system. Let  $x$  denote the generic vector of unknowns (the dimension of  $x$  will depend on which of the three systems is to be solved). We may, for example, solve  $Py = b$ , where  $P$  is one of the three appropriate preconditioners, or we may set  $y_1 = 0$ ,  $y_2 = D^{-1} b_2$ . Then use PCG with preconditioner  $P$  to solve

$$B\hat{x} = b - By \quad \text{and set} \quad x = y + \hat{x}.$$

In general, if  $b_2 \neq 0$ , it is not safe to apply the PCG method to the indefinite system (1.2). Moreover, the PCG method will usually generate different iterates for the condensed system (1.3) and the doubly augmented system (1.4).

Finally, we note that the condensed system (1.3) and doubly augmented system (1.4) may be viewed as being preconditioned versions of each other, as defined in the following result.

PROPOSITION 3.2. *Consider the PCG method applied to a generic symmetric system  $Ax = b$  with symmetric positive-definite preconditioner  $P$  and initial iterate  $x_0 = 0$ . Let  $L$  be a nonsingular matrix with the same dimension as  $A$ . Then, if the PCG method is applied to  $LAL^T \hat{x} = Lb$  with preconditioner  $LPL^T$  and initial iterate  $\hat{x}_0 = 0$ , the PCG iterates are related by the transformation  $x = L^T \hat{x}$ .*

If we consider the decomposition

$$\begin{pmatrix} M + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix} = \begin{pmatrix} I & A^T D^{-1} \\ & I \end{pmatrix} \begin{pmatrix} M + A^T D^{-1} A & \\ & D \end{pmatrix} \begin{pmatrix} I & \\ D^{-1} A & I \end{pmatrix},$$

then Proposition 3.2 implies that the doubly augmented system may be viewed as a particular preconditioned version of the condensed system augmented by the diagonal  $D$  for the  $x_2$  variables (or vice versa). This is a further illustration that the proposed approach gives equal weight to  $x_1$  and  $x_2$ . We prefer to do the analysis in terms of the doubly augmented system because it provides the parameterization based on the scalar parameter  $\nu$ .

**3.1. Properties of the constraint preconditioners.** We now embed the preconditioners of (3.1) within a family of preconditioners, parameterized by the scalar  $\nu$ . This parameterization is analogous to the parameterization of the matrices of Proposition 2.2. The parameterization allows a unified analysis of the three preconditioners given in (3.1).

DEFINITION 3.3 (a parameterized preconditioner). *Associated with the matrix  $B(\nu)$  of Definition 2.1, we define the preconditioner  $P(\nu)$  as*

$$P(\nu) = \begin{pmatrix} M + (1 + \nu)A^T D^{-1} A & \nu A^T \\ \nu A & \nu D \end{pmatrix},$$

where  $M$  is a symmetric approximation to  $H$  such that

- (P<sub>1</sub>)  $\|M\| = O(1)$ ;
- (P<sub>2</sub>)  $M + A^T D^{-1} A$  is positive definite;
- (P<sub>3</sub>)  $\text{eig}_{\min}(M + A^T D^{-1} A) = \Omega(1)$ .

Given suitable  $A$  and  $D$ , a matrix  $M$  satisfying the conditions of Definition 3.3 may be found, for example, by using a suitable factorization when solving with  $P(-1)$  (see Forsgren and Murray [17], Forsgren and Gill [16], and Forsgren [18]).

Proposition 2.3 gives  $\text{In}(P(\nu)) = \text{In}(M + A^T D^{-1} A) + \text{In}(\nu D)$ . It follows that  $P(\nu)$  is nonsingular for  $\nu \neq 0$ , positive definite for  $\nu > 0$ , and  $\text{In}(P(\nu)) = (n, m, 0)$  for  $\nu < 0$ . It is straightforward to show that for all nonzero  $\nu$ , the eigenvalues of  $P(\nu)^{-1} B(\nu)$  are real and independent of  $\nu$ . The first lemma reveals the structure of  $P(\nu)^{-1} B(\nu)$ .

LEMMA 3.4 (structure of the parameterized preconditioner). *Let  $B(\nu)$  and  $P(\nu)$  be defined as in Definitions 2.1 and 3.3, respectively. Then, for  $\nu \neq 0$ , it holds that*

$$(3.2) \quad P(\nu)^{-1} B(\nu) = \begin{pmatrix} S & \\ T & I \end{pmatrix} = \begin{pmatrix} I & \\ -D^{-1} A & I \end{pmatrix} \begin{pmatrix} S & \\ & I \end{pmatrix} \begin{pmatrix} I & \\ D^{-1} A & I \end{pmatrix},$$

where  $S$  and  $T$  are given by

$$(3.3a) \quad S = (M + A^T D^{-1} A)^{-1} (H + A^T D^{-1} A),$$

$$(3.3b) \quad T = D^{-1} A (M + A^T D^{-1} A)^{-1} (M - H).$$

In addition, the spectrum of  $P(\nu)^{-1} B(\nu)$  is independent of  $\nu$  and consists of  $m$  unit eigenvalues and the  $n$  eigenvalues of  $(M + A^T D^{-1} A)^{-1} (H + A^T D^{-1} A)$ .

*Proof.* The expressions for  $P(\nu)^{-1} B(\nu)$  follow from the decomposition given in Lemma A.1 of the appendix. The similarity transform (3.2) implies that the spectrum of  $P(\nu)^{-1} B(\nu)$  consists of the  $n$  eigenvalues of  $(M + A^T D^{-1} A)^{-1} (H + A^T D^{-1} A)$  together with  $m$  unit eigenvalues.  $\square$

Next we relate the  $O(\mu)$  diagonal elements in  $D$  to eigenvalues of size  $1 + O(\mu^{1/2})$  in the (1, 1) block  $S$  of  $P(\nu)^{-1} B(\nu)$  from (3.2).

LEMMA 3.5 (eigenvalues of the parameterized preconditioner). *Let  $M$  satisfy assumptions (P<sub>1</sub>)–(P<sub>3</sub>) of Definition 3.3. Let  $A_{\mathcal{S}}$  denote the submatrix of rows of  $A$  associated with diagonal elements of  $D$  that are  $O(\mu)$ . Then the eigenvalues of*

$$(M + A^T D^{-1} A)^{-1} (H + A^T D^{-1} A)$$

are all  $O(1)$  with at least  $\text{rank}(A_{\mathcal{S}})$  being  $1 + O(\mu^{1/2})$ .

*Proof.* First we show that  $(M + A^T D^{-1} A)^{-1}$  has at least  $\text{rank}(A_{\mathcal{S}})$  eigenvalues that are  $O(\mu)$ . Let  $m_1 = \text{rank}(A_{\mathcal{S}})$ . Without loss of generality it may be assumed that the rows of  $A$  are ordered so that the  $m_1$  row indices in  $\mathcal{S}$  corresponding to linearly



independent rows of  $A_S$  appear first. This implies that  $A$  and  $D$  may be partitioned as

$$A = \begin{pmatrix} A_1 \\ A_2 \end{pmatrix} \quad \text{and} \quad D = \begin{pmatrix} D_1 & \\ & D_2 \end{pmatrix},$$

where  $A_1$  is  $m_1 \times n$ , and  $D_1$  is  $m_1 \times m_1$  with all eigenvalues  $O(\mu)$ . Then

$$A^T D^{-1} A = A_1^T D_1^{-1} A_1 + A_2^T D_2^{-1} A_2.$$

Consider the singular-value decomposition  $A_1 = U \Sigma V^T$ , where  $U$  and  $V$  are orthonormal matrices of dimension  $m_1 \times m_1$  and  $n \times m_1$ , respectively, and  $\Sigma$  is an  $m_1 \times m_1$  diagonal matrix. Let  $v = Vp$ , where  $p$  is an arbitrary  $m_1$ -vector of unit length. Then

$$(3.4) \quad v^T A_1^T D_1^{-1} A_1 v = p^T V^T A_1^T D_1^{-1} A_1 V p = p^T \Sigma U^T D_1^{-1} U \Sigma p \geq \sigma_{m_1}^2 \text{eig}_{\min}(D_1^{-1}).$$

Assumption (A<sub>3</sub>) implies that  $\sigma_{m_1} = \Theta(1)$ . In addition, all eigenvalues of  $D_1$  are  $O(\mu)$ , and so (3.4) implies that  $v^T A_1^T D_1^{-1} A_1 v = \Omega(1/\mu)$ . It follows that

$$(3.5) \quad v^T (M + A^T D^{-1} A) v = v^T M v + v^T A_1^T D_1^{-1} A_1 v + v^T A_2^T D_2^{-1} A_2 v = \Omega(1/\mu),$$

since  $v^T M v = O(1)$  and  $v^T A_2^T D_2^{-1} A_2 v \geq 0$ . As  $p$  is an arbitrary unit vector such that  $v = Vp$ , we conclude from (3.5) that there exists an  $m_1$ -dimensional subspace of vectors  $v$  such that  $v^T (M + A^T D^{-1} A) v = \Omega(1/\mu)$ . The Courant–Fischer min-max theorem implies that  $M + A^T D^{-1} A$  has at least  $m_1$  eigenvalues that are  $\Omega(1/\mu)$ ; see, e.g., [28, Theorem 8.1.2, p. 394]. It follows that there exists an  $n \times m_1$  orthonormal matrix  $Y$  and  $m_1 \times m_1$  diagonal  $\Lambda$  with  $\text{eig}_{\min}(\Lambda) = \Omega(1/\mu)$  such that  $(M + A^T D^{-1} A)Y = Y\Lambda$ . Since  $M + A^T D^{-1} A$  is positive definite, it must hold that  $(M + A^T D^{-1} A)^{-1} Y = Y\Lambda^{-1}$  and  $(M + A^T D^{-1} A)^{-1/2} Y = Y\Lambda^{-1/2}$  with

$$(3.6a) \quad \|(Y^T (M + A^T D^{-1} A)^{-1} Y)\| = O(\mu),$$

$$(3.6b) \quad \|(M + A^T D^{-1} A)^{-1/2} Y\| = O(\mu^{1/2}).$$

Let  $E = (M + A^T D^{-1} A)^{-1/2} (H - M) (M + A^T D^{-1} A)^{-1/2}$  and let  $Z$  be an  $n \times (n - m_1)$  orthonormal matrix such that the columns of  $Z$  form a basis for the null space of  $Y^T$ . Then  $Q = (Y \ Z)$  is orthonormal and  $E$  has the same eigenvalues as the matrix

$$Q^T E Q = \begin{pmatrix} Y^T E Y & Y^T E Z \\ Z^T E Y & Z^T E Z \end{pmatrix},$$

where

$$(3.7a) \quad Y^T E Y = Y^T (M + A^T D^{-1} A)^{-1/2} (H - M) (M + A^T D^{-1} A)^{-1/2} Y,$$

$$(3.7b) \quad Y^T E Z = Y^T (M + A^T D^{-1} A)^{-1/2} (H - M) (M + A^T D^{-1} A)^{-1/2} Z,$$

$$(3.7c) \quad Z^T E Z = Z^T (M + A^T D^{-1} A)^{-1/2} (H - M) (M + A^T D^{-1} A)^{-1/2} Z.$$

Then Definition 3.3 and the order estimates (3.6) imply that  $\|Y^T E Y\| = O(\mu)$ ,  $\|Z^T E Y\| = O(\mu^{1/2})$ , and  $\|Z^T E Z\| = O(1)$ . Hence, since  $\|Z^T E Y\| = O(\mu^{1/2})$ , the eigenvalues of  $E$  differ by  $O(\mu^{1/2})$  from the eigenvalues of  $Y^T E Y$  together with the

eigenvalues of  $Z^T E Z$ . But, since  $\|Y^T E Y\| = O(\mu)$ , we conclude that  $E$  has  $m_1$  eigenvalues that are  $O(\mu^{1/2})$ . By similarity,

$$\text{eig}((M + A^T D^{-1} A)^{-1}(H - M)) = \text{eig}(E),$$

and it must hold that  $(M + A^T D^{-1} A)^{-1}(H - M)$  has at least  $m_1$  eigenvalues that are  $O(\mu^{1/2})$ . The required results now follow from the identity

$$(M + A^T D^{-1} A)^{-1}(H + A^T D^{-1} A) = I + (M + A^T D^{-1} A)^{-1}(H - M),$$

completing the proof.  $\square$

If  $M - H$  is known to be a definite matrix, then the  $O(\mu^{1/2})$  bound of Lemma 3.5 may be sharpened to be  $O(\mu)$ . In this case, the  $O(\mu)$  curvature of the product  $(M + A^T D^{-1} A)^{-1/2}(H - M)(M + A^T D^{-1} A)^{-1/2}$  over a  $\text{rank}(A_S)$ -dimensional space implied by (3.7a) is sufficient to guarantee  $\text{rank}(A_S)$  eigenvalues  $1 + O(\mu)$ .

A combination of Lemmas 3.4 and 3.5 gives the following result on the eigenvalues of  $P(\nu)^{-1}B(\nu)$ .

**THEOREM 3.6** (eigenvalues of the preconditioned matrix). *Let  $\nu \neq 0$ , and let  $B(\nu)$  and  $P(\nu)$  be defined as in Definitions 2.1 and 3.3, respectively. Let  $A_S$  denote the submatrix of rows of  $A$  associated with diagonal elements of  $D$  that are  $O(\mu)$ . The preconditioned matrix  $P(\nu)^{-1}B(\nu)$  has the following properties:*

- (a) *The spectrum of  $P(\nu)^{-1}B(\nu)$  is independent of  $\nu$  and consists of  $m$  unit eigenvalues and the  $n$  eigenvalues of  $(M + A^T D^{-1} A)^{-1}(H + A^T D^{-1} A)$ .*
- (b) *Every eigenvalue of  $P(\nu)^{-1}B(\nu)$  is of order  $O(1)$ . Moreover,  $P(\nu)^{-1}B(\nu)$  has at least  $m + \text{rank}(A_S)$  eigenvalues  $1 + O(\mu^{1/2})$ , of which at least  $m$  are exactly one.*

This result implies that if  $m_S$  denotes the number of eigenvalues of  $D$  that are  $O(\mu)$  and the corresponding  $m_S \times n$  submatrix  $A_S$  has full row rank, then the PCG method can be expected to give a solution that is  $O(\mu^{1/2})$  accurate in at most  $n - m_S$  iterations.

**4. Active-set preconditioning.** An advantage of interior methods is that all inequality constraints are treated in the same way—i.e., the solution path does not depend on an explicit prediction of which constraints are active at the solution. However, this advantage also can be a weakness because all constraint gradients are included in the linear system, even those having little or no influence on the solution. For example, if an interior method is applied to a problem with 100 variables and 100,000 inequality constraints, then a KKT system with 100,100 rows and columns must be solved at each iteration. However, if only 50 (say) of the inequalities are active at the solution, an active-set method would need to solve a KKT system of order 150. In the context of an interior method, the partition of constraints into “active” and “inactive” is determined by the magnitude of the diagonals of  $D$  in the KKT system (1.2). Broadly speaking, the active set at the solution is estimated by the indices of the “small” diagonals, and the inactive set is estimated by the indices of the “big” diagonals.

In this section we formulate and analyze two *active-set* preconditioners based on discarding rows of  $A$  that correspond to the big diagonals of  $D$ . The preconditioners may be applied with a cost comparable to that of solving the KKT system in an active-set method. In addition, the preconditioners allow considerable flexibility in how the diagonals are partitioned into large and small elements—the partition affects

only the rate of convergence of the iterative solver, not the rate of convergence of the interior method. Similar preconditioners have been proposed by Gertz and Griffin [21] in the context of support vector machine classifiers for large data sets. Preconditioners for the solution of linear programs in standard form have been considered by Gill et al. [25] and Oliveira and Sorensen [38]. An active-set preconditioner for general nonlinear optimization has been proposed by Lukšan, Matonoha, and Vlček [35].

**4.1. Two active-set preconditioners.** Let  $m_{\mathcal{S}}$ ,  $m_{\mathcal{M}}$ , and  $m_{\mathcal{B}}$  denote the number of row indices in the sets  $\mathcal{S}$ ,  $\mathcal{M}$ , and  $\mathcal{B}$  of “small,” “medium,” and “big” elements of  $D$  (see section 1.3). These sets are disjoint, and together they contain all the row indices of  $A$ , so that  $m_{\mathcal{S}} + m_{\mathcal{M}} + m_{\mathcal{B}} = m$ . If strict complementarity holds for the underlying optimization problem, then  $m_{\mathcal{M}}$  is zero for all  $\mu$  sufficiently small (see, e.g., Forsgren, Gill, and Wright [15, p. 531]). The following analysis, does not assume strict complementarity and so  $m_{\mathcal{M}}$  may be nonzero. However, it must be emphasized that in this situation, the assumption regarding the order of the small elements of  $D$  is a simplification of the real situation. Our assumption that  $d_{ii} = O(\mu)$  for  $i \in \mathcal{S}$  is sufficient to capture the behavior as  $\mu$  converges to zero. For a detailed discussion regarding interior methods on degenerate problems, see, e.g., Wright and Orban [48].

In order to simplify the notation, the indices corresponding to the small and medium diagonals are combined into one set  $\mathcal{C}$ , i.e.,  $\mathcal{C} = \mathcal{S} \cup \mathcal{M}$ . This set is the complement of  $\mathcal{B}$ , i.e.,  $\mathcal{C} \cap \mathcal{B} = \emptyset$  and  $\mathcal{C} \cup \mathcal{B} = \{1, \dots, m\}$ . This simplification is possible because of our focus on preconditioners based on discarding information associated with the indices in  $\mathcal{B}$ . Given the partition induced by  $\mathcal{B}$  and  $\mathcal{C}$ , the matrix  $P(\nu)$  may be partitioned as

$$(4.1) \quad P(\nu) = \begin{pmatrix} M + (1 + \nu)A^T D^{-1} A & \nu A_{\mathcal{C}}^T & \nu A_{\mathcal{B}}^T \\ \nu A_{\mathcal{C}} & \nu D_{\mathcal{C}} & \\ \nu A_{\mathcal{B}} & & \nu D_{\mathcal{B}} \end{pmatrix}.$$

By eliminating the  $\nu D_{\mathcal{B}}$  block from  $P(\nu)$ , we may factor  $P(\nu)$  as  $P(\nu) = R_{\mathcal{P}} P_{\mathcal{P}}(\nu) R_{\mathcal{P}}^T$ , with

$$(4.2a) \quad R_{\mathcal{P}} = \begin{pmatrix} I & & A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} \\ & I & \\ & & I \end{pmatrix},$$

$$(4.2b) \quad P_{\mathcal{P}}(\nu) = \begin{pmatrix} M + A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}} + (1 + \nu)A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}} & \nu A_{\mathcal{C}}^T & \\ \nu A_{\mathcal{C}} & \nu D_{\mathcal{C}} & \\ & & \nu D_{\mathcal{B}} \end{pmatrix}.$$

Here, the subscript “ $\mathcal{P}$ ” identifies matrices that depend on the partition induced by  $\mathcal{B}$  and  $\mathcal{C}$  (note that  $P(\nu)$  itself is independent of the partition).

The nontrivial step associated with applying the preconditioner in the factored form  $P(\nu) = R_{\mathcal{P}} P_{\mathcal{P}}(\nu) R_{\mathcal{P}}^T$  requires a solve with the leading principal submatrix of (4.2b):

$$(4.3) \quad \begin{pmatrix} M + A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}} + (1 + \nu)A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}} & \nu A_{\mathcal{C}}^T \\ \nu A_{\mathcal{C}} & \nu D_{\mathcal{C}} \end{pmatrix}.$$

This matrix, formed by eliminating the block  $\nu D_{\mathcal{B}}$  from  $P(\nu)$ , has smaller dimension,  $(n + m - m_{\mathcal{B}}) \times (n + m - m_{\mathcal{B}})$ , compared to  $(n + m) \times (n + m)$  for  $P(\nu)$ . Lukšan, Matonoha, and Vlček [35] propose an active-set preconditioner based on forming an

incomplete factorization of the  $(1, 1)$  block. This avoids unnecessary fill-in from the term  $A_B^T D_B^{-1} A_B$ . We propose an alternative strategy based on the observation that since  $\|D_B^{-1}\| = O(\mu)$ , then  $\|A_B^T D_B^{-1} A_B\| = O(\mu)$ . In particular, the term  $A_B^T D_B^{-1} A_B$  may be omitted from the  $(1, 1)$  block of  $P_{\mathcal{P}}(\nu)$  without significantly changing the preconditioner. This implies that (4.3) is replaced by  $P_{\mathcal{C}}(\nu)$ , where

$$(4.4) \quad P_{\mathcal{C}}(\nu) = \begin{pmatrix} M + (1 + \nu)A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}} & \nu A_{\mathcal{C}}^T \\ \nu A_{\mathcal{C}} & \nu D_{\mathcal{C}} \end{pmatrix}.$$

The subscript “ $\mathcal{C}$ ” indicates that  $P_{\mathcal{C}}(\nu)$  depends only on the indices in  $\mathcal{C}$ . In active-set constraint preconditioning,  $P_{\mathcal{C}}(\nu)$  plays the role of  $P(\nu)$  in the analysis of the standard case in section 3.1. Analogous to the assumptions on  $M$ ,  $A$ , and  $D$  in Definition 3.3, we require that

- (P'<sub>1</sub>)  $\|M\| = O(1)$ ;
- (P'<sub>2</sub>)  $M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}}$  is positive definite; and
- (P'<sub>3</sub>)  $\text{eig}_{\min}(M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}}) = \Omega(1)$ .

When  $P_{\mathcal{C}}(\nu)$  replaces the leading principal submatrix in (4.2b) the product of the factors becomes

$$(4.5) \quad P_{\mathcal{P}}^1(\nu) = \begin{pmatrix} M + \nu A_B^T D_B^{-1} A_B + (1 + \nu)A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}} & \nu A_{\mathcal{C}}^T & \nu A_B^T \\ \nu A_{\mathcal{C}} & \nu D_{\mathcal{C}} & \\ \nu A_B & & \nu D_B \end{pmatrix},$$

which alternatively may be viewed as the preconditioner obtained by subtracting the term  $A_B^T D_B^{-1} A_B$  from the  $(1, 1)$  block of  $P(\nu)$ .

The preconditioner (4.5) has the factorization  $P_{\mathcal{P}}^1(\nu) = R_{\mathcal{P}} P_{\mathcal{P}}^2(\nu) R_{\mathcal{P}}^T$ , where  $R_{\mathcal{P}}$  is the upper-triangular matrix (4.2a) and  $P_{\mathcal{P}}^2(\nu)$  is given by

$$(4.6) \quad P_{\mathcal{P}}^2(\nu) = \begin{pmatrix} M + (1 + \nu)A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}} & \nu A_{\mathcal{C}}^T & \\ \nu A_{\mathcal{C}} & \nu D_{\mathcal{C}} & \\ & & \nu D_B \end{pmatrix}.$$

The matrix  $P_{\mathcal{P}}^2(\nu)$  is yet another active-set preconditioner, which may be derived differently by replacing the leading principal submatrix of  $P_{\mathcal{P}}(\nu)$  by  $P_{\mathcal{C}}(\nu)$  and replacing  $R_{\mathcal{P}}$  by  $I$ . Observe that the replacement of  $R_{\mathcal{P}}$  by  $I$  quantifies the difference between  $P_{\mathcal{P}}^1(\nu)$  and  $P_{\mathcal{P}}^2(\nu)$ , and hence  $P_{\mathcal{P}}^1(\nu)$  is always a “better” approximation to  $P_{\mathcal{P}}(\nu)$  than  $P_{\mathcal{P}}^2(\nu)$ . However, regardless of the choice of preconditioner, the dominant cost is the solve with the matrix  $P_{\mathcal{C}}(\nu)$  of (4.4). Note that  $A_B$  does not appear in  $P_{\mathcal{P}}^2(\nu)$ , which may make  $P_{\mathcal{P}}^2(\nu)$  the more attractive preconditioner when it is expensive to form  $A$ , e.g., in PDE-constrained optimization [5].

It remains to establish the theoretical properties of the preconditioners  $P_{\mathcal{P}}^1(\nu)$  and  $P_{\mathcal{P}}^2(\nu)$ . The next result shows that, asymptotically, the eigenvalues of  $P(\nu)^{-1}B(\nu)$  and  $P_{\mathcal{P}}^1(\nu)^{-1}B(\nu)$  are identical.

**THEOREM 4.1** (properties of the preconditioner  $P_{\mathcal{P}}^1(\nu)$ ). *Let  $B(\nu)$  and  $P_{\mathcal{P}}^1(\nu)$  be as defined in Definition 2.1 and (4.5), respectively. In addition, assume that assumptions (P'<sub>1</sub>)–(P'<sub>3</sub>) hold. Then  $P_{\mathcal{P}}^1(\nu)$  is positive definite for all  $\nu > 0$ . Moreover, the following properties hold for all  $\nu \neq 0$ :*

- (a) *The spectrum of  $P_{\mathcal{P}}^1(\nu)^{-1}B(\nu)$  is independent of  $\nu$  and consists of  $m$  unit eigenvalues and the  $n$  eigenvalues of  $(M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1}(H + A^T D^{-1} A)$ .*
- (b) *The matrix  $P_{\mathcal{P}}^1(\nu)^{-1}B(\nu)$  has all eigenvalues of order  $O(1)$  and at least  $m + \text{rank}(A_S)$  eigenvalues  $1 + O(\mu^{1/2})$ , of which at least  $m$  are exactly one.*

(c) If  $P(\nu)$  is the preconditioner of Definition 3.3, then  $\text{eig}(P_{\mathcal{P}}^1(\nu)^{-1}B(\nu)) = \text{eig}(P(\nu)^{-1}B(\nu)) + O(\mu)$ .

*Proof.* Proposition 2.3 implies that the preconditioner  $P_{\mathcal{P}}^1(\nu)$  is positive definite for all  $\nu > 0$ . For the remainder of the proof it will be assumed that  $\nu$  is nonzero. It follows that  $\text{eig}(P_{\mathcal{P}}^1(\nu)^{-1}B(\nu)) = \text{eig}(P_{\mathcal{P}}^2(\nu)^{-1}R_{\mathcal{P}}^{-1}B(\nu)R_{\mathcal{P}}^{-T})$ , where

$$(4.7) \quad R_{\mathcal{P}}^{-1}B(\nu)R_{\mathcal{P}}^{-T} = \begin{pmatrix} H + A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}} + (1 + \nu) A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}} & \nu A_{\mathcal{C}}^T & \\ & \nu A_{\mathcal{C}} & \nu D_{\mathcal{C}} \\ & & \nu D_{\mathcal{B}} \end{pmatrix}.$$

By successively replacing  $H$  by  $H + A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}}$ ,  $A$  by  $A_{\mathcal{C}}$ , and  $D$  by  $D_{\mathcal{C}}$  in Lemma 3.4, a combination of (4.6) and (4.7) gives

$$(4.8) \quad P_{\mathcal{P}}^2(\nu)^{-1}R_{\mathcal{P}}^{-1}B(\nu)R_{\mathcal{P}}^{-T} = \begin{pmatrix} S_{\mathcal{C}} & & \\ T_{\mathcal{C}} & I & \\ & & I \end{pmatrix},$$

where the matrices  $S_{\mathcal{C}}$  and  $T_{\mathcal{C}}$  are given by

$$(4.9a) \quad S_{\mathcal{C}} = (M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1} (H + A^T D^{-1} A),$$

$$(4.9b) \quad T_{\mathcal{C}} = D_{\mathcal{C}}^{-1} A_{\mathcal{C}} (M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1} (M - H - A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}}).$$

The identity (4.8) implies that the spectrum of  $P_{\mathcal{P}}^1(\nu)^{-1}B(\nu)$  consists of the eigenvalues of  $S_{\mathcal{C}}$  and  $m$  unit eigenvalues, which proves part (a). Since  $S_{\mathcal{C}}$  is independent of  $\nu$ , the spectrum of  $P_{\mathcal{P}}^1(\nu)^{-1}B(\nu)$  is also independent of  $\nu$ . Lemma 3.5 implies that  $S_{\mathcal{C}}$  has at least  $\text{rank}(A_{\mathcal{S}})$  eigenvalues that are  $1 + O(\mu^{1/2})$ , which establishes part (b).

To establish part (c), we need to estimate the difference between the eigenvalues of  $S_{\mathcal{C}}$  and  $S$ , where  $S$  is given by (3.3a). This can be done using Lemma A.2 of the appendix. We may write

$$(4.10a) \quad S_{\mathcal{C}} = I + (M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1} (H - M) + (M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1} A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}},$$

$$(4.10b) \quad S = I + (M + A^T D^{-1} A)^{-1} (H - M).$$

By assumption, matrix  $M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}}$  is positive definite with the smallest eigenvalue bounded away from zero, and  $A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}}$  is positive definite with  $\|D_{\mathcal{B}}^{-1}\| = O(\mu)$ . The identity (4.10a) implies that the matrix  $(M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{1/2} S_{\mathcal{C}} (M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1/2}$  is symmetric and has the same eigenvalues as  $S_{\mathcal{C}}$ , and it follows from (4.10) that

$$(4.11a) \quad \text{eig}(S_{\mathcal{C}}) = 1 + \text{eig}((M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1} (H - M)) + O(\mu),$$

$$(4.11b) \quad \text{eig}(S) = 1 + \text{eig}((M + A^T D^{-1} A)^{-1} (H - M)).$$

If we define  $M_1 = M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}}$ ,  $M_2 = A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}}$ , and  $M_3 = H - M$ , then (4.11) gives  $\text{eig}(S_{\mathcal{C}}) = 1 + \text{eig}(M_1^{-1} M_3) + O(\mu)$  and  $\text{eig}(S) = 1 + \text{eig}((M_1 + M_2)^{-1} M_3)$ . Lemma A.2 in conjunction with assumptions (P'<sub>1</sub>)–(P'<sub>3</sub>) gives the desired result.  $\square$

Next we establish that  $P_{\mathcal{P}}^2(\nu)$  has the same asymptotic behavior as  $P(\nu)$  and  $P_{\mathcal{P}}^1(\nu)$ . For  $P_{\mathcal{P}}^2(\nu)$  it is assumed that  $\nu > 0$ , which ensures that the eigenvalues of  $P_{\mathcal{P}}^2(\nu)^{-1}B(\nu)$  are real. The preconditioner  $P_{\mathcal{P}}^2(\nu)$  is less expensive to apply than  $P_{\mathcal{P}}^1(\nu)$ , but the number of unit eigenvalues of the preconditioned matrix decreases from  $m$  to  $m - \text{rank}(A_{\mathcal{B}})$  because  $A_{\mathcal{B}}$  does not appear in  $P_{\mathcal{P}}^2(\nu)$ . However, as the next theorem shows, for  $\nu > 0$ ,  $P_{\mathcal{P}}^2(\nu)$  behaves almost as well as  $P_{\mathcal{P}}^1(\nu)$  in the sense that the eigenvalues of  $P_{\mathcal{P}}^2(\nu)^{-1}B(\nu)$  differ from the eigenvalues of  $P(\nu)^{-1}B(\nu)$  by  $O(\mu^{1/2})$ .

THEOREM 4.2 (properties of the preconditioner  $P_{\mathcal{P}}^2(\nu)$ ). *Let  $B(\nu)$  and  $P_{\mathcal{P}}^2(\nu)$  be as defined in Definition 2.1 and (4.6), respectively. In addition, assume that assumptions  $(P'_1)$ – $(P'_3)$  hold. Then the following properties hold for all  $\nu > 0$ :*

- (a)  $P_{\mathcal{P}}^2(\nu)$  is positive definite.
- (b) The matrix  $P_{\mathcal{P}}^2(\nu)^{-1}B(\nu)$  has all eigenvalues of order  $O(1)$  and at least  $m + \text{rank}(A_{\mathcal{S}})$  eigenvalues  $1 + O(\mu^{1/2})$ , of which at least  $m - \text{rank}(A_{\mathcal{B}})$  are exactly one.
- (c) If  $P(\nu)$  is the preconditioner of Definition 3.3, then  $\text{eig}(P_{\mathcal{P}}^2(\nu)^{-1}B(\nu)) = \text{eig}(P(\nu)^{-1}B(\nu)) + O(\mu^{1/2})$ .

*Proof.* The positive definiteness of  $P_{\mathcal{P}}^2(\nu)$  for  $\nu > 0$  follows from Proposition 2.3. For the remainder of the proof, assume that  $\nu > 0$ . Then, since  $P_{\mathcal{P}}^2(\nu)$  is positive definite, the identity  $\text{eig}(P_{\mathcal{P}}^2(\nu)^{-1/2}B(\nu)P_{\mathcal{P}}^2(\nu)^{-1/2}) = \text{eig}(P_{\mathcal{P}}^2(\nu)^{-1}B(\nu))$  ensures that  $P_{\mathcal{P}}^2(\nu)^{-1}B(\nu)$  has real eigenvalues. Analogous to the proof of Theorem 4.1, by successively replacing the matrix  $H$  by  $H + (1 + \nu)A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}}$ ,  $A$  by  $A_{\mathcal{C}}$ , and  $D$  by  $D_{\mathcal{C}}$  in Lemma 3.4, and by using a combination of Proposition 2.3 and (4.6), we find that

$$(4.12) \quad P_{\mathcal{P}}^2(\nu)^{-1}B(\nu) = \begin{pmatrix} S_{\mathcal{C}} & & \\ T_{\mathcal{C}} & I & \\ & & I \end{pmatrix} + \begin{pmatrix} U & X \\ V & Y \\ W & \end{pmatrix},$$

where  $S_{\mathcal{C}}$  and  $T_{\mathcal{C}}$  are given by (4.9),

- (4.13a)  $U = \nu(M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1} A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}},$
- (4.13b)  $V = -\nu D_{\mathcal{C}}^{-1} A_{\mathcal{C}} (M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1} A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}},$
- (4.13c)  $W = D_{\mathcal{B}}^{-1} A_{\mathcal{B}},$
- (4.13d)  $X = \nu(M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1} A_{\mathcal{B}}^T,$
- (4.13e)  $Y = -\nu D_{\mathcal{C}}^{-1} A_{\mathcal{C}} (M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}})^{-1} A_{\mathcal{B}}^T.$

It follows from (4.12) that  $P_{\mathcal{P}}^2(\nu)^{-1}B(\nu)$  contains  $m - m_{\mathcal{B}}$  columns from the identity matrix; hence it has at least  $m - m_{\mathcal{B}}$  unit eigenvalues. The remaining eigenvalues are those of the matrix  $N$  given by

$$(4.14) \quad N = \begin{pmatrix} S_{\mathcal{C}} + U & X \\ W & I \end{pmatrix} = \begin{pmatrix} S_{\mathcal{C}} + \nu S_M^{-1} A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}} & \nu S_M^{-1} A_{\mathcal{B}}^T \\ D_{\mathcal{B}}^{-1} A_{\mathcal{B}} & I \end{pmatrix},$$

with  $S_M = M + A_{\mathcal{C}}^T D_{\mathcal{C}}^{-1} A_{\mathcal{C}}$ . Observe that (4.14) implies that any nonzero vector  $x$  such that  $A_{\mathcal{B}}^T x = 0$  induces an eigenvector corresponding to a unit eigenvalue of  $N$ . Hence,  $N$  has at least  $m_{\mathcal{B}} - \text{rank}(A_{\mathcal{B}})$  unit eigenvalues. Further, let  $\tilde{Q}$  be defined by

$$\tilde{Q} = \begin{pmatrix} S_M^{1/2} & \\ & \nu^{1/2} D_{\mathcal{B}}^{1/2} \end{pmatrix}.$$

Then  $N$  and  $\tilde{Q}N\tilde{Q}^{-1}$  have identical eigenvalues, and  $\tilde{Q}N\tilde{Q}^{-1}$  is given by

$$\begin{aligned} \tilde{Q}N\tilde{Q}^{-1} &= \begin{pmatrix} S_M^{1/2} S_{\mathcal{C}} S_M^{-1/2} + \nu S_M^{-1/2} A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}} S_M^{-1/2} & \nu^{1/2} S_M^{-1/2} A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1/2} \\ \nu^{1/2} D_{\mathcal{B}}^{-1/2} A_{\mathcal{B}} S_M^{-1/2} & I \end{pmatrix} \\ &= \begin{pmatrix} S_M^{1/2} S_{\mathcal{C}} S_M^{-1/2} & \\ & I \end{pmatrix} + \begin{pmatrix} \nu S_M^{-1/2} A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}} S_M^{-1/2} & \nu^{1/2} S_M^{-1/2} A_{\mathcal{B}}^T D_{\mathcal{B}}^{-1/2} \\ \nu^{1/2} D_{\mathcal{B}}^{-1/2} A_{\mathcal{B}} S_M^{-1/2} & \end{pmatrix}. \end{aligned}$$

Note that  $S_M^{1/2} S_C S_M^{-1/2}$  is symmetric, and hence  $\tilde{Q} N \tilde{Q}^{-1}$  is symmetric. In addition, from our assumptions, it follows that  $\|\nu^{1/2} S_M^{-1/2} A_B^T D_B^{-1/2}\| = O(\mu^{1/2})$  and  $\|\nu S_M^{-1/2} A_B^T D_B^{-1} A_B S_M^{-1/2}\| = O(\mu)$ . Hence,  $\tilde{Q} N \tilde{Q}^{-1}$  has  $m_B$  eigenvalues that differ by  $O(\mu^{1/2})$  from unity, and  $n$  eigenvalues that differ by  $O(\mu^{1/2})$  from the eigenvalues of the matrix  $S_M^{1/2} S_C S_M^{-1/2}$ . In addition, the eigenvalues of  $S_M^{1/2} S_C S_M^{-1/2}$  and  $S_C$  are identical. Consequently, it follows that the spectrum of  $P_{\mathcal{P}}^2(\nu)^{-1} B(\nu)$  consists of  $m - \text{rank}(A_B)$  unit eigenvalues,  $\text{rank}(A_B)$  eigenvalues that are  $1 + O(\mu^{1/2})$ , and  $n$  eigenvalues that differ by  $O(\mu^{1/2})$  from the eigenvalues of  $S_C$ . Theorem 4.1 now shows that  $\text{eig}(P_{\mathcal{P}}^2(\nu)^{-1} B(\nu)) = \text{eig}(P_{\mathcal{P}}^1(\nu)^{-1} B(\nu)) + O(\mu^{1/2})$ , which gives the required result, since  $O(\mu^{1/2})$  dominates  $O(\mu)$ . In particular, Theorem 4.1 implies that  $P_{\mathcal{P}}^2(\nu)^{-1} B(\nu)$  has at least  $m + \text{rank}(A_S)$  eigenvalues that are  $1 + O(\mu^{1/2})$ .  $\square$

We conclude that it is possible to construct appropriate constraint preconditioners based on solving the smaller system (4.4). Moreover, the matrix  $P_C(\nu)$  of (4.4) has exactly the same structure as  $P(\nu)$ . The difference is that the number of rows and columns in the preconditioner has been reduced from  $n + m$  to  $n + m - m_B$ . Hence, all the previous analysis applies. For our example with 100 variables and 100,000 inequality constraints, a matrix of dimension 150 would need to be factored instead of a matrix of dimension 100,100.

As shown above, the partition of the row indices into  $\mathcal{B}$  and its complement  $\mathcal{C}$  provides active-set preconditioners  $P_{\mathcal{P}}^1(\nu)$  and  $P_{\mathcal{P}}^2(\nu)$  that are asymptotically equivalent to  $P(\nu)$ . If strict complementarity holds, then  $m_{\mathcal{M}} = 0$  and the division into large and small elements is straightforward. If strict complementarity does not hold, then  $m_{\mathcal{M}} > 0$  and we have chosen to append  $\mathcal{M}$  to  $\mathcal{S}$ . Analogous preconditioners may be constructed by first identifying  $\mathcal{S}$  and then forming the complementary set  $\bar{\mathcal{S}}$ , which is the set obtained by appending  $\mathcal{M}$  to  $\mathcal{B}$ . The resulting KKT system analogous to (4.4) would have smaller dimension because  $\mathcal{C}$  is replaced by  $\mathcal{S}$ . However, the resulting preconditioners would not be asymptotically equivalent in general. For  $P_{\mathcal{P}}^1(\nu)$ , the  $1 + O(\mu^{1/2})$  cluster of eigenvalues would be the same as for  $P(\nu)$ , but the eigenvalues resulting from  $M$  would differ by an  $O(1)$  term. The reason for this difference is that the norm of  $D_{\bar{\mathcal{S}}}^{-1}$  would not be of order  $O(\mu)$ , but would include terms involving  $m_{\mathcal{M}}$  eigenvalues of order one.

It should be emphasized that the choice of  $\mathcal{C}$  and  $\mathcal{B}$  affects only the efficiency of the active-set constraint preconditioners and not the definition of the linear equations that need to be solved. A poorly chosen partition may adversely affect the quality of the preconditioner, but not the solution of the linear equations. The partition analyzed here provides the largest  $\mathcal{B}$  for which we can guarantee that the preconditioners  $P_{\mathcal{P}}^1(\nu)$  and  $P_{\mathcal{P}}^2(\nu)$  are asymptotically equivalent to  $P(\nu)$  for  $\nu > 0$ . If elements are excluded from  $\mathcal{B}$ , then  $P_{\mathcal{P}}^1(\nu)$  and  $P_{\mathcal{P}}^2(\nu)$  become “better” approximations to  $P(\nu)$ , and the asymptotic performance is unchanged. However, this increases the dimension of the KKT system (4.4). As noted above, if  $\mathcal{B}$  is chosen too large, in the sense that diagonal elements of  $D$  are included in  $D_{\mathcal{B}}$  that are not  $\Omega(1/\mu)$ , then the quality of the active-set preconditioners can be expected to deteriorate. Hence, it is not essential that  $\mathcal{B}$  is estimated correctly, but it is essential that  $D_{\mathcal{B}}$  contains only large elements.

**5. On semidefinite diagonal matrices.** Up to this point we have assumed that the matrix  $D$  in the (2, 2) block of the KKT system is positive definite. In the general case, the last block of equations in the KKT system has the form

$$(5.1) \quad Ax_1 + Gx_2 = b_2,$$

where  $G$  is a diagonal matrix with positive and zero entries. If all the constraints of the optimization problem are nonlinear, it is always possible to formulate the interior method so that  $G$  is positive definite. For inequality constraints, standard formulations give positive elements in  $G$  that are of the order of the perturbation parameter  $\mu$  (see, e.g., Vanderbei and Carpenter [46] and Forsgren and Gill [16]). Typically, zero elements of  $G$  are associated with linearized equality constraints, where the corresponding subset of equations (5.1) are the Newton equations for a zero of the constraint residual. An alternative to direct constraint linearization is to impose equality constraints approximately via a quadratic penalty function. It can be shown that this approach gives a positive element in  $G$  of the order of  $\bar{\mu}$ , where  $\bar{\mu}$  is the inverse of the penalty parameter (see, e.g., Gould [31] and Forsgren and Gill [16]). The parameter  $\bar{\mu}$  may be allowed to vary with  $\mu$ , or may be fixed at some small value (see, e.g., Gill et al. [26] and Saunders and Tomlin [42]). Fixing  $\bar{\mu}$  defines a *regularization* of the problem, which allows the formulation of methods that do not require an assumption on the rank of the equality constraint Jacobian. (For more details on the use of regularization in interior methods, see Gill et al. [24], Vanderbei and Shanno [47], and Altman and Gondzio [1].)

However, it may not always be beneficial to regularize *linear* constraints. Regularization in this context is less crucial because reliable techniques exist for discarding dependent equality constraints. Moreover, interior methods can be defined so that every iterate satisfies the linear equality constraints (see below). With an appropriate choice of constraints, this feature can be used to guarantee that the nonlinear functions and their derivatives are well defined at all points generated by the interior method.

In order to consider KKT systems with a semidefinite (2,2) block, we assume that the variables and equations are preordered to give a system  $Bx = b$  such that

$$(5.2) \quad \begin{pmatrix} H & -A^T & -F^T \\ -A & -D & \\ -F & & \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} b_1 \\ -b_2 \\ -b_3 \end{pmatrix},$$

where  $D$  is positive definite. Note that we cannot compute the condensed or doubly augmented system for these equation because of the zero block. In this case,  $B$  has correct inertia if  $N^T(H + A^T D^{-1} A)N$  is positive definite, where the columns of  $N$  form a basis for the null space of  $F$  (see Forsgren [18]).

The KKT system (5.2) may be solved using a projection technique similar to that described in section 3. First, an initial point  $y$  is found with first  $n$  components forming a vector  $y_1$  such that  $Fy_1 = b_3$ . This vector may be computed in various ways—e.g., by computing an LU factorization of  $F^T$  (see, e.g., Gill, Murray, and Saunders [27]), or by solving a system for the preconditioning matrix associated with (5.2), where  $H$  is replaced by a suitable approximation  $M$  (see Gould, Hribar, and Nocedal [29]). Once  $y$  is known, the PCG method may be used to solve

$$(5.3) \quad \begin{pmatrix} H & -A^T & -F^T \\ -A & -D & \\ -F & & \end{pmatrix} \begin{pmatrix} \hat{x}_1 \\ \hat{x}_2 \\ \hat{x}_3 \end{pmatrix} = \begin{pmatrix} \hat{b}_1 \\ -\hat{b}_2 \\ 0 \end{pmatrix},$$

with  $\hat{b}_1 = b_1 - Hy_1 + A^T y_2 + F^T y_3$  and  $\hat{b}_2 = b_2 - Ay_1 - Dy_2$ . The required solution is then  $x = y + \hat{x}$ . Analogous to the situation when  $G$  is positive definite, we may



embed (5.3) into the parameterized system of linear equations

$$(5.4) \quad \begin{pmatrix} H + (1 + \nu)A^T D^{-1}A & \nu A^T & -F^T \\ \nu A & \nu D & \\ -F & & \end{pmatrix} \begin{pmatrix} \hat{x}_1 \\ \hat{x}_2 \\ \hat{x}_3 \end{pmatrix} = \begin{pmatrix} \hat{b}_1 + (1 + \nu)A^T D^{-1}\hat{b}_2 \\ \nu \hat{b}_2 \\ 0 \end{pmatrix},$$

with  $\hat{b}_1 = b_1 - Hy_1 + A^T y_2 + F^T y_3$  and  $\hat{b}_2 = b_2 - Ay_1 - Dy_2$ . If the zero elements of  $G$  are associated with linear constraints, and the system (5.3) is solved exactly, it suffices to compute the special step  $y$  only once, when solving the first system. Then, provided that the constraint preconditioner is applied exactly at every PCG step, the right-hand side of (5.3) will remain zero for all subsequent iterations.

The linear equations (5.2), (5.3), and (5.4) do not require  $N$ . If the preconditioner cannot be applied exactly, then it is necessary to use an alternative method based on computing products of the form  $N^T v$  and  $Nu$ . (Gill, Murray, and Saunders [27] describe how these products may be computed in a numerically stable way without needing to store  $N$  explicitly.) The requirement that  $F\hat{x}_1 = 0$  implies that  $\hat{x}_1$  can be written as  $\hat{x}_1 = N\hat{p}_1$ . Substituting this expression in (5.3) gives the reduced KKT system

$$(5.5) \quad \begin{pmatrix} N^T H N & -(AN)^T \\ -AN & -D \end{pmatrix} \begin{pmatrix} \hat{p}_1 \\ \hat{p}_2 \end{pmatrix} = \begin{pmatrix} N^T(b_1 - Hy_1 + A^T y_2) \\ -b_2 + Ay_1 + Dy_2 \end{pmatrix},$$

from which we can define  $\hat{x}_1 = N\hat{p}_1$  and  $\hat{x}_2 = \hat{p}_2$ . This system has a nonsingular  $(2, 2)$  block and has correct inertia if (5.2) has correct inertia. Moreover, the iterates define exact projections regardless of the accuracy of the solves with the preconditioner. Hence, all the conditions needed for the application of the PCG method proposed in section 3 apply. The systems (5.3) and (5.5) are mathematically equivalent, which implies that we may apply the analysis of section 3 directly to both (5.3) and (5.4).

**6. Some numerical examples.** To illustrate the numerical performance of the proposed preconditioners, a PCG method was applied to a collection of illustrative large sparse KKT systems. The test matrices were generated from a number of realistic KKT systems arising in the context of primal-dual interior methods. We conclude with some randomly generated problems that illustrate some of the properties of the preconditioned matrices.

**6.1. Examples from the COPS test set.** First we describe some numerical results obtained on linear equations arising in a primal-dual interior method applied to optimization problems from the COPS 3.0 test collection [6, 9, 10] implemented in the AMPL modeling language [2, 19].

The equations are analogous to those generated by an interior-point method with barrier parameter  $\mu$ . The data for the test matrices was generated using a primal-dual trust-region method (see, e.g., [16, 20, 32]) applied to eight problems, *Camshape*, *Channel*, *Gasoil*, *Marine*, *Methanol*, *Pinene*, *Polygon*, and *Tetra*, from the COPS 3.0 test collection [6, 8, 9, 10]. The interior-point method requires the solution of systems with a KKT matrix of the form

$$(6.1) \quad \begin{pmatrix} H & -J^T \\ -J & -\Gamma \end{pmatrix},$$

where  $H$  is the  $n \times n$  Hessian of the Lagrangian,  $J$  is the  $m \times n$  Jacobian matrix of constraint gradients, and  $\Gamma$  is a positive-definite diagonal with some large and small

TABLE 6.1  
 Dimensions of the AMPL versions of the COPS problems.

Problem	$n$	$m$	$\ H - M\ $	$\sigma_k(J)$
<i>Camshape</i>	1200	3600	2.6e+0	1.3e-5
<i>Channel</i>	6398	6398	1.1e+2	4.1e-5
<i>Gasoil</i>	4001	4001	1.1e+1	0.0e+0
<i>Marine</i>	6415	6407	3.5e+1	0.0e+0
<i>Methanol</i>	4802	4802	3.1e+0	2.8e-3
<i>Pinene</i>	8000	8000	9.2e+3	6.8e-8
<i>Polygon</i>	398	20496	2.4e+2	0.0e+0
<i>Tetra</i>	1200	4254	2.7e+1	0.0e+0

elements. These systems have the same structure as the generic system (1.2). The dimensions of the eight problems are given in Table 6.1. The optimization problems in the COPS collection have a mixture of general nonlinear constraints and simple upper and lower bounds on the variables. The simple bounds lead to unit rows in  $J$ , and it is customary to define a smaller KKT system in which the unit rows and columns are eliminated. However, in the numerical experiments, the unit rows were included in order to more accurately illustrate the results of Theorems 3.6, 4.1, and 4.2. Hence the value of  $m$  also includes the bound constraints. The final column gives the  $k$ th largest singular value of  $J$ , where  $k = \min\{m, n\}$ .

For each of the eight featured COPS problems, matrices  $H$ ,  $A$ ,  $D$  and the right-hand side were generated from the matrices  $H$ ,  $J$ ,  $\Gamma$  and the right-hand side at a snapshot taken at iteration 30 of the interior method. For each problem snapshot, five systems of equations were generated by specifying five matrices  $D$  with entries parameterized by a scalar  $\mu = 10^{-\ell}$  for  $\ell = \{1, 2, 4, 6, 8\}$ . For each value of  $\mu$ , the matrices  $A$  and  $D$  were generated from  $J$  and  $\Gamma$  using the MATLAB code fragment

```
[D, ind] = sort(Gamma); % sort the diagonals of Gamma
A = J(ind,:);          % reorder the rows of J
k = min([m,n]);       %
if k < m, D(k+1:end) = max(D(k+1:end), 1/mu); end
D(1:k) = min(D(1:k),mu);
```

This choice of  $D$  implicitly defines a sequence of systems associated with a vertex solution of the underlying optimization problem for which strict complementarity holds. This was done deliberately to minimize the effect of the matrix  $M$  on the efficiency of the preconditioner (see the definition of  $P(\nu)$  in (4.1)). Asymptotically, the matrix  $M$  defines the efficiency of the preconditioner within the null space of the matrix of active constraints (see, e.g., Dollar et al. [11]). In our analysis we have focused on the part of the preconditioner associated with the constraint part of the KKT system. The formulation and analysis of effective choices for  $M$  are beyond the scope of this paper. (For some possible approaches, see, e.g., [22, 23].) In the experiments reported here,  $M$  was a diagonal matrix with entries  $M_{jj} = \max(|H_{jj}|, \delta)$ , where  $\delta = 10^{-1}$ .

Figure 6.1 depicts the number of PCG iterations required to solve the resulting 8 sets of 5 systems of linear equations. The bar charts give the PCG iterations for the condensed system (top) and doubly augmented system (bottom). The MATLAB version of SYMMLQ [39] was used as the PCG solver. The symmetric indefinite solver MA27 was used to factor the constraint preconditioner (see Duff and Reid [13]). The

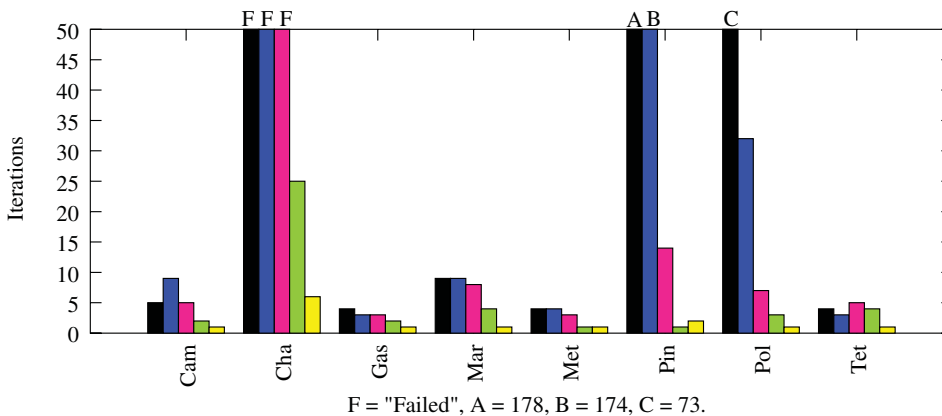
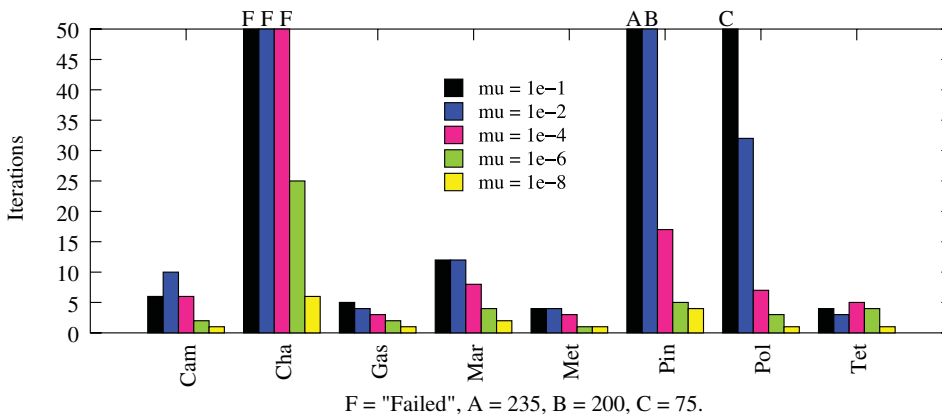


FIG. 6.1. Results on the COPS problems. The graphs give the number of PCG iterations for  $P(0)$  (above) and  $P(1)$  (below). The vertical axis is limited to 50 iterations. The number of iterations needed for each off-scale case is given by the appropriate key code "A", "B", or "C".

value of  $10^{-6}$  was used for the SYMMLQ relative convergence tolerance. For problem *Channel*, with the three larger values of  $\mu$ , PCG did not converge within the pre-assigned limit of  $10^6$  iterations (a more sophisticated choice of  $M$  is needed in this case). Note the similar number of PCG iterations needed to solve the condensed system and doubly augmented system.

Figure 6.2 gives the number of PCG iterations for the active-set preconditioners on the COPS problems *Camshape*, *Polygon*, and *Tetra*. These problems have significantly more constraints than variables and provide good examples on which to test the active-set preconditioners. In order to illustrate the behavior of the active-set preconditioner, we scale  $D$  so that exactly  $n$  elements are less than  $\mu$  and the remaining elements are greater than  $1/\mu$ . We emphasize that the motivation for manipulating  $D$  in this way is to illustrate the effect of changing  $\mu$  for fixed  $H$  and  $J$ .

```

mS = min([n,m,k]);
JS = J(1:mS,:);    JB = J(mS+1:end,:);
DS = D(1:mS);      DB = D(mS+1:end);
DS = spdiags(DS,0,mS,mS);

```

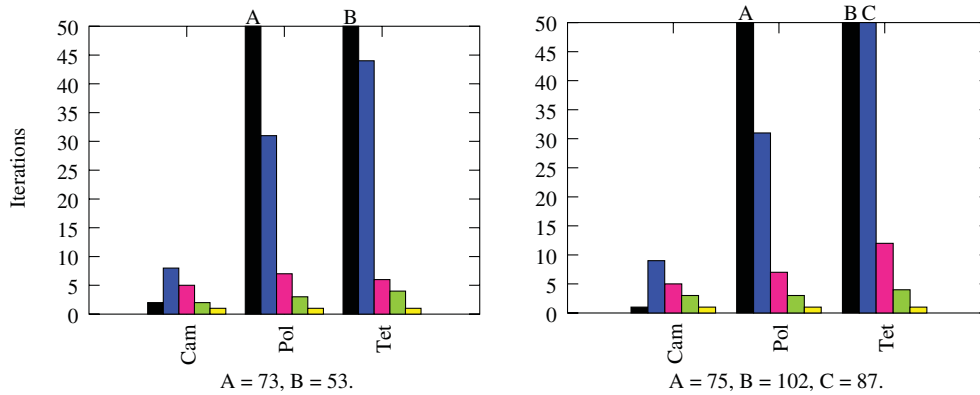


FIG. 6.2. *COPS* problems: PCG iterations for  $P_p^1(1)$  (left) and  $P_p^2(1)$  (right). The vertical axis is limited to 50 iterations. The number of iterations needed for each off-scale case is given by the appropriate key code “A”, “B”, or “C”.

**6.2. Results from randomly generated problems.** Additional experiments were performed on randomly generated KKT systems. The purpose of these experiments was to illustrate the clustering of the eigenvalues of the preconditioned matrices associated with the doubly augmented system. The first set of experiments involved applying the preconditioners  $P(1)$ ,  $P_p^1(1)$ , and  $P_p^2(1)$  to randomly generated problems satisfying the assumptions of Theorems 4.1 and 4.2. Of particular interest is the “strict complementarity” assumption that every element of the diagonal  $D$  is either big or small. Given values  $n = 400$ ,  $m = 600$ , and  $\mu = 10^{-\ell}$  for  $\ell = \{1, 2, 4, 6, 8\}$ , matrices  $H$ ,  $A_S$ ,  $A_B$ ,  $D_S$ , and  $D_B$  were generated using the MATLAB code fragment

```
mS = 100; d = 10^(-2);
H = sprandsym(n,d);      JS = sprand(mS,n,d,0.1);  JB = sprand(m-mS,n,d);
DS = diag(mu*ones(mS,1));  DB = diag((1/mu)*ones(m-mS,1));
```

Table 6.2 gives details of the eigenvalues of the preconditioned matrices associated with each of the preconditioners  $P(1)$ ,  $P_p^1(1)$ , and  $P_p^2(1)$ , where the diagonal preconditioner  $M$  was defined as in the *COPS* examples of the previous section. In all these runs, the resulting KKT matrix satisfies  $\|H - M\| = 4.95$  and  $\sigma_n(J_S) = 10^{-1}$ . These linear systems would be typical for a primal-dual method applied to an optimization problem with 100 active constraints at a point satisfying a strict complementarity assumption. Theorems 3.6 and 4.1 predict that for the preconditioners  $P(1)$  and  $P_p^1(1)$ , 700 ( $= m + \text{rank}(A_S)$ ) eigenvalues of the preconditioned matrix will cluster close to unity, with 600 of these eigenvalues exactly equal to one. Theorem 4.2 predicts that as  $\mu$  is reduced,  $P_p^2(1)$  also will give 700 eigenvalues close to one, whereas 200 ( $= m - n$ ) eigenvalues will be exactly one.

The last four columns of Table 6.2 illustrate the degree of clustering of the eigenvalues of the preconditioned matrix. Clustering is measured by means of the function  $l(\theta)$  defined as follows. Given a matrix  $C$  with real eigenvalues, the function

$$l(\theta) = \text{card}\{\lambda \in \text{eig}(C) : |\lambda - 1| \leq \theta\}$$

gives the number of eigenvalues of  $C$  within distance  $\theta$  of unity. Table 6.2 gives the values of  $l(\theta)$  for the three preconditioned matrices  $C = P(1)^{-1}B(1)$ ,  $P_p^1(1)^{-1}B(1)$ , and  $P_p^2(1)^{-1}B(1)$ . In this strict-complementarity case, we expect that the proposed preconditioners would asymptotically give a cluster of 700 unit eigenvalues. Note that

TABLE 6.2

Number of clustered eigenvalues of the preconditioned matrix. Randomly generated KKT systems with  $n = 400$ ,  $m = 600$ ,  $m_S = 100$ ,  $m_M = 0$ ,  $m_B = 500$ ,  $(D_S)_{ii} = \mu$ , and  $(D_B)_{ii} = 1/\mu$ .

	$\mu$	$l(10^{-8})$	$l(10^{-6})$	$l(10^{-4})$	$l(10^{-2})$
$P$	$10^{-1}$	600	600	600	601
	$10^{-2}$	600	600	600	612
	$10^{-4}$	600	600	612	675
	$10^{-6}$	600	612	673	700
	$10^{-8}$	606	673	700	700
$P_P^1$	$10^{-1}$	600	600	600	602
	$10^{-2}$	600	600	600	611
	$10^{-4}$	600	600	612	675
	$10^{-6}$	600	612	673	700
	$10^{-8}$	593	673	700	700
$P_P^2$	$10^{-1}$	200	200	200	425
	$10^{-2}$	200	200	215	567
	$10^{-4}$	200	215	566	673
	$10^{-6}$	215	566	672	700
	$10^{-8}$	546	672	700	700

TABLE 6.3

Number of clustered eigenvalues for the preconditioned matrix. Randomly generated KKT systems with  $n = 400$ ,  $m = 600$ ,  $m_S = 75$ ,  $m_M = 25$ ,  $m_B = 500$ ,  $(D_S)_{ii} = \mu$ ,  $(D_M)_{ii} = 1$ , and  $(D_B)_{ii} = 1/\mu$ .

	$\mu$	$l(10^{-8})$	$l(10^{-6})$	$l(10^{-4})$	$l(10^{-2})$
$P$	$10^{-1}$	600	600	600	601
	$10^{-2}$	600	600	600	609
	$10^{-4}$	600	600	609	654
	$10^{-6}$	600	609	653	675
	$10^{-8}$	604	653	675	675
$P_P^1$	$10^{-1}$	600	600	600	602
	$10^{-2}$	600	600	600	609
	$10^{-4}$	600	600	609	654
	$10^{-6}$	600	609	653	675
	$10^{-8}$	591	653	675	675
$P_P^2$	$10^{-1}$	200	200	200	425
	$10^{-2}$	200	200	215	563
	$10^{-4}$	200	215	563	653
	$10^{-6}$	215	563	653	675
	$10^{-8}$	544	653	675	675

for small values of  $\mu$ ,  $P(1)$  and  $P_P^1(1)$  produce very similar numbers of eigenvalues close to unity. The preconditioner  $P_P^2(1)$  tends to give fewer accurate eigenvalues than  $P(1)$  and  $P_P^1(1)$  for the larger values of  $\mu$ , although the differences become less marked as  $\mu$  is reduced.

Table 6.3 was generated with the same data used for Table 6.2, with the one exception that strict complementarity was assumed not to hold. As in Table 6.2, we simulate an optimization problem with 100 active constraints, but in this case we set  $m_S = 75$  and  $m_M = 25$ . The corresponding diagonal elements of  $D_M$  were set at one.

In this non-strict-complementarity case, Theorems 3.6, 4.1, and 4.2 predict that the proposed preconditioners would asymptotically give a cluster of 675 ( $= m + m_S$ ) unit eigenvalues, which is reflected in the results. The performance of the preconditioners is very similar to that depicted in Table 6.2.

**7. Summary and further research.** A framework has been proposed for applying the PCG method to KKT systems of the form (1.1) that arise in interior methods for general nonconvex optimization. The proposed methods are based on applying the conjugate-gradient method to the doubly augmented system (1.4), which is positive definite if the underlying optimization problem satisfies the second-order sufficient conditions for optimality. An advantage of the doubly augmented system is that it is positive definite with respect to all of the variables.

We also have proposed a class of constraint preconditioners for the doubly augmented system. In particular, we have analyzed two ways of using an estimate of the active set to reduce the cost of applying the preconditioner when there are many inequality constraints. As the solution of the optimization problem is approached, these active-set preconditioners have theoretical performance comparable to constraint preconditioners that include all the constraints. An advantage of using preconditioning in conjunction with the doubly augmented system is that the linear equations used to apply the preconditioner need not be solved exactly. Future work will consider the analysis associated with these approximate preconditioners.

The focus of this paper has been on the formulation and analysis of *constraint* preconditioners. The next step is to consider “full” preconditioners based on estimating the matrix  $H$  in the  $(1, 1)$  block of the KKT equations. For example, a preconditioner may be defined using an incomplete inertia-controlling factorization of the KKT system (1.2). For more details on the inertia-controlling factorization for augmented systems in interior methods, see Forsgren and Gill [16] and Forsgren [18].

**Appendix. Linear algebra.** Here we review two results from linear algebra. The first gives the structure of the inverse of  $B(\nu)$  and may be verified by direct multiplication.

LEMMA A.1. *Given a nonsingular symmetric matrix  $D$ , consider the matrix*

$$B(\nu) = \begin{pmatrix} H + (1 + \nu)A^T D^{-1} A & \nu A^T \\ \nu A & \nu D \end{pmatrix},$$

where  $\nu$  is a scalar. Then  $B(\nu)$  may be factored in the form

$$B(\nu) = \begin{pmatrix} I & A^T D^{-1} \\ & I \end{pmatrix} \begin{pmatrix} H + A^T D^{-1} A & \\ & \nu D \end{pmatrix} \begin{pmatrix} I & \\ D^{-1} A & I \end{pmatrix}.$$

Moreover, if  $H + A^T D^{-1} A$  is nonsingular and  $\nu \neq 0$ , then  $B(\nu)$  is nonsingular, with inverse

$$\begin{aligned} B(\nu)^{-1} &= \begin{pmatrix} (H + A^T D^{-1} A)^{-1} & -(H + A^T D^{-1} A)^{-1} A^T D^{-1} \\ -D^{-1} A (H + A^T D^{-1} A)^{-1} & \frac{1}{\nu} D^{-1} + D^{-1} A (H + A^T D^{-1} A)^{-1} A^T D^{-1} \end{pmatrix} \\ &= \begin{pmatrix} I & \\ -D^{-1} A & I \end{pmatrix} \begin{pmatrix} (H + A^T D^{-1} A)^{-1} & \\ & \frac{1}{\nu} D^{-1} \end{pmatrix} \begin{pmatrix} I & -A^T D^{-1} \\ & I \end{pmatrix}. \end{aligned}$$

The second result provides bounds on the perturbation of the eigenvalues of  $M_1^{-1} M_3$  when  $M_1$  is perturbed by a positive-semidefinite matrix  $M_2$ .

LEMMA A.2. Let  $M_1$ ,  $M_2$ , and  $M_3$  be  $n \times n$  symmetric matrices with  $M_1$  positive definite and  $M_2$  positive semidefinite. Let  $\{\lambda_i\}$  and  $\{\tilde{\lambda}_i\}$  denote the eigenvalues of  $M_1^{-1}M_3$  and  $(M_1 + M_2)^{-1}M_3$ , respectively. Assume that the  $\{\lambda_i\}$  are ordered so that  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ , with the same ordering for  $\{\tilde{\lambda}_i\}$ . Then

$$\begin{aligned} 0 &\leq \lambda_i - \tilde{\lambda}_i \leq \|M_1^{-1/2}M_2M_1^{-1/2}\|\tilde{\lambda}_i \quad \text{for all } i \text{ such that } \lambda_i \geq 0, \\ 0 &\leq \tilde{\lambda}_i - \lambda_i \leq -\|M_1^{-1/2}M_2M_1^{-1/2}\|\tilde{\lambda}_i \quad \text{for all } i \text{ such that } \lambda_i < 0. \end{aligned}$$

*Proof.* Let  $M$ ,  $\tilde{M}$ , and  $\tilde{I}$  be defined such that

$$M = M_1^{-1/2}M_3M_1^{-1/2}, \quad \tilde{I} = I + M_1^{-1/2}M_2M_1^{-1/2}, \quad \text{and} \quad \tilde{M} = \tilde{I}^{-1/2}M\tilde{I}^{-1/2}.$$

Then  $M$ ,  $\tilde{I}$ , and  $\tilde{M}$  are symmetric with  $\tilde{I}$  positive definite. A similarity transformation gives  $\text{eig}(M_1^{-1}M_3) = \text{eig}(M_1^{-1/2}M_3M_1^{-1/2}) = \text{eig}(M)$ , which means that  $M$  has eigenvalues  $\lambda_i$ ,  $i = 1:n$ . Similarly, we have

$$\begin{aligned} (M_1 + M_2)^{-1}M_3 &= M_1^{-1/2}(I + M_1^{-1/2}M_2M_1^{-1/2})^{-1}M_1^{-1/2}M_3 \\ \text{(A.1)} \quad &= M_1^{-1/2}\tilde{I}^{-1}MM_1^{1/2}. \end{aligned}$$

Successive similarity transformations of (A.1) with  $M_1^{1/2}$  and  $\tilde{I}^{1/2}$  give

$$\text{eig}((M_1 + M_2)^{-1}M_3) = \text{eig}(\tilde{I}^{-1}M) = \text{eig}(\tilde{I}^{-1/2}M\tilde{I}^{-1/2}) = \text{eig}(\tilde{M}),$$

which means that  $\tilde{M}$  has eigenvalues  $\tilde{\lambda}_i$ ,  $i = 1:n$ .

Now we relate  $\lambda_i$  to  $\tilde{\lambda}_i$ . First, consider the case  $\lambda_i \geq 0$ . Since  $\lambda_i$  is an eigenvalue of  $M$  and  $\tilde{\lambda}_i$  is an eigenvalue of  $\tilde{M}$ , with  $\tilde{M} = \tilde{I}^{-1/2}M\tilde{I}^{-1/2}$ , the Courant–Fischer min-max theorem gives

$$\text{(A.2)} \quad \frac{1}{\|\tilde{I}^{-1}\|}\tilde{\lambda}_i \leq \lambda_i \leq \|\tilde{I}\|\tilde{\lambda}_i;$$

see, e.g., Golub and Van Loan [28, pp. 403–404]. Since  $\tilde{I} = I + M_1^{-1/2}M_2M_1^{-1/2}$  with  $M_2$  positive semidefinite, it follows that  $\|\tilde{I}^{-1}\| \leq 1$  and  $\|\tilde{I}\| \leq 1 + \|M_1^{-1/2}M_2M_1^{-1/2}\|$ . Hence, (A.2) gives

$$\tilde{\lambda}_i \leq \lambda_i \leq (1 + \|M_1^{-1/2}M_2M_1^{-1/2}\|)\tilde{\lambda}_i,$$

which is equivalent to the desired result when  $\lambda_i \geq 0$ ,

$$0 \leq \lambda_i - \tilde{\lambda}_i \leq \|M_1^{-1/2}M_2M_1^{-1/2}\|\tilde{\lambda}_i.$$

For the case  $\lambda_i < 0$ , we apply the analysis above to the matrices  $-\tilde{M}$  and  $-M$ . Then, since  $-\lambda_i$  is a positive eigenvalue of  $-M$  and  $-\tilde{\lambda}_i$  is an eigenvalue of  $-\tilde{M}$ , we conclude that

$$0 \leq -\lambda_i + \tilde{\lambda}_i \leq -\|M_1^{-1/2}M_2M_1^{-1/2}\|\tilde{\lambda}_i,$$

as required.  $\square$

**Acknowledgments.** The authors would like to thank Michael Gertz for many stimulating discussions on the theoretical and practical aspects of interior methods. The authors are also grateful to the referees for many constructive comments that significantly improved the presentation.

## REFERENCES

- [1] A. ALTMAN AND J. GONDZIO, *Regularized symmetric indefinite systems in interior point methods for linear and quadratic optimization*, Optim. Methods Softw., 11/12 (1999), pp. 275–302.
- [2] AMPL HOME PAGE, <http://www.ampl.com>.
- [3] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, in Acta Numerica, 2005, Acta Numer. 14, Cambridge University Press, Cambridge, UK, 2005, pp. 1–137.
- [4] L. BERGAMASCHI, J. GONDZIO, AND G. ZILLI, *Preconditioning indefinite systems in interior point methods for optimization*, Comput. Optim. Appl., 28 (2004), pp. 149–171.
- [5] G. BIROS AND O. GHATTAS, *Inexactness issues in the Lagrange-Newton-Krylov-Schur method for PDE-constrained optimization*, in Large-Scale PDE-Constrained Optimization (Santa Fe, NM, 2001), Lect. Notes Comput. Sci. Eng. 30, Springer, Berlin, 2003, pp. 93–114.
- [6] A. BONDARENKO, D. BORTZ, AND J. J. MORÉ, *COPS: Large-Scale Nonlinearly Constrained Optimization Problems*, Technical Report ANL/MCS-TM-237, Mathematics and Computer Science division, Argonne National Laboratory, Argonne, IL, 1998 (revised October 1999).
- [7] R. S. DEMBO, S. C. EISENSTAT, AND T. STEIHAUG, *Inexact Newton methods*, SIAM J. Numer. Anal., 19 (1982), pp. 400–408.
- [8] E. D. DOLAN, J. J. MORÉ, AND T. S. MUNSON, *Benchmarking Optimization Software with COPS 3.0*, Technical Memorandum ANL/MCS-TM-273, Argonne National Laboratory, Argonne, IL, 2004.
- [9] E. D. DOLAN AND J. J. MORÉ, *Benchmarking Optimization Software with COPS*, Technical Memorandum ANL/MCS-TM-246, Argonne National Laboratory, Argonne, IL, 2000.
- [10] E. D. DOLAN AND J. J. MORÉ, *Benchmarking optimization software with performance profiles*, Math. Program., 91 (2002), pp. 201–213.
- [11] H. S. DOLLAR, N. I. M. GOULD, W. H. A. SCHILDERS, AND A. J. WATHEN, *On Iterative Methods and Implicit-Factorization Preconditioners for Regularized Saddle-Point Systems*, Report RAL-TR-2005-011, Rutherford Appleton Laboratory, Oxfordshire, UK, 2005.
- [12] H. S. DOLLAR, *Extending Constraint Preconditioners for Saddle Point Problems*, Numerical Analysis Group Research Report NA-05/02, Oxford University Computing Laboratory, Oxford, UK, 2005.
- [13] I. S. DUFF AND J. K. REID, *MA27: A Set of Fortran Subroutines for Solving Sparse Symmetric Sets of Linear Equations*, Technical Report R-10533, Computer Science and Systems Division, AERE Harwell, Oxford, UK, 1982.
- [14] H. C. ELMAN, V. E. HOWLE, J. N. SHADID, AND R. S. TUMINARO, *A parallel block multi-level preconditioner for the 3d incompressible Navier-Stokes equations*, J. Comput. Phys., 187 (2003), pp. 504–523.
- [15] A. FORSGREN, P. E. GILL, AND M. H. WRIGHT, *Interior methods for nonlinear optimization*, SIAM Rev., 44 (2002), pp. 525–597.
- [16] A. FORSGREN AND P. E. GILL, *Primal-dual interior methods for nonconvex nonlinear programming*, SIAM J. Optim., 8 (1998), pp. 1132–1152.
- [17] A. FORSGREN AND W. MURRAY, *Newton methods for large-scale linear equality-constrained minimization*, SIAM J. Matrix Anal. Appl., 14 (1993), pp. 560–587.
- [18] A. FORSGREN, *Inertia-controlling factorizations for optimization algorithms*, Appl. Numer. Math., 43 (2002), pp. 91–107.
- [19] R. FOURER, D. M. GAY, AND B. W. KERNIGHAN, *AMPL: A Modeling Language for Mathematical Programming*, Brooks/Cole-Thomson Learning, Pacific Grove, CA, 2003.
- [20] E. M. GERTZ AND P. E. GILL, *A primal-dual trust-region algorithm for nonlinear programming*, Math. Program. Ser. B, 100 (2004), pp. 49–94.
- [21] E. M. GERTZ AND J. D. GRIFFIN, *Support Vector Machine Classifiers for Large Data Sets*, Technical Memorandum ANL/MCS-TM-289, Mathematics and Computer Science division, Argonne National Laboratory, Argonne, IL, 2005.
- [22] O. GHATTAS AND J.-H. BARK, *Large-scale SQP methods for optimization of Navier-Stokes flows*, in Large-Scale Optimization with Applications, Part II (Minneapolis, MN, 1995), Springer, New York, 1997, pp. 247–270.



- [23] O. GHATTAS AND C. E. OROZCO, *A parallel reduced Hessian SQP method for shape optimization*, in Multidisciplinary Design Optimization (Hampton, VA, 1995), SIAM, Philadelphia, PA, 1997, pp. 133–152.
- [24] P. E. GILL, W. MURRAY, D. B. PONCELEÓN, AND M. A. SAUNDERS, *Primal-Dual Methods for Linear Programming*, Report SOL 91-3, Department of Operations Research, Stanford University, Stanford, CA, 1991.
- [25] P. E. GILL, W. MURRAY, D. B. PONCELEÓN, AND M. A. SAUNDERS, *Preconditioners for indefinite systems arising in optimization*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 292–311.
- [26] P. E. GILL, W. MURRAY, D. B. PONCELEÓN, AND M. A. SAUNDERS, *Solving reduced KKT systems in barrier methods for linear programming*, in Numerical Analysis 1993 (Dundee, 1993), G. A. Watson and D. Griffiths, eds., Pitman Res. Notes Math. Ser. 303, Longman Sci. Tech., Harlow, UK, 1994, pp. 89–104.
- [27] P. E. GILL, W. MURRAY, AND M. A. SAUNDERS, *SNOPT: An SQP algorithm for large-scale constrained optimization*, SIAM Rev., 47 (2005), pp. 99–131.
- [28] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.
- [29] N. I. M. GOULD, M. E. HRIBAR, AND J. NOCEDAL, *On the solution of equality constrained quadratic programming problems arising in optimization*, SIAM J. Sci. Comput., 23 (2001), pp. 1376–1395.
- [30] N. I. M. GOULD, S. LUCIDI, M. ROMA, AND PH. L. TOINT, *Solving the trust-region subproblem using the Lanczos method*, SIAM J. Optim., 9 (1999), pp. 504–525.
- [31] N. I. M. GOULD, *On the accurate determination of search directions for simple differentiable penalty functions*, IMA J. Numer. Anal., 6 (1986), pp. 357–372.
- [32] J. D. GRIFFIN, *Interior-Point Methods for Large-Scale Nonconvex Optimization*, Ph.D. thesis, Department of Mathematics, University of California, San Diego, CA, 2005.
- [33] W. W. HAGER, *Minimizing a quadratic over a sphere*, SIAM J. Optim., 12 (2001), pp. 188–208.
- [34] C. KELLER, N. I. M. GOULD, AND A. J. WATHEN, *Constraint preconditioning for indefinite linear systems*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1300–1317.
- [35] L. LUKŠAN, C. MATONOHA, AND J. VLČEK, *Interior-point method for non-linear non-convex optimization*, Numer. Linear Algebra Appl., 11 (2004), pp. 431–453.
- [36] L. LUKŠAN AND J. VLČEK, *Indefinitely preconditioned inexact Newton method for large sparse equality constrained non-linear programming problems*, Numer. Linear Algebra Appl., 5 (1998), pp. 219–247.
- [37] Y. NOTAY, *Flexible conjugate gradients*, SIAM J. Sci. Comput., 22 (2000), pp. 1444–1460.
- [38] A. R. L. OLIVEIRA AND D. C. SORESENSEN, *A new class of preconditioners for large-scale linear systems from interior point methods for linear programming*, Linear Algebra Appl., 394 (2005), pp. 1–24.
- [39] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
- [40] I. PERUGIA AND V. SIMONCINI, *Block-diagonal and indefinite symmetric preconditioners for mixed finite element formulations*, Numer. Linear Algebra Appl., 7 (2000), pp. 585–616.
- [41] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM J. Sci. Comput., 14 (1993), pp. 461–469.
- [42] M. A. SAUNDERS AND J. A. TOMLIN, *Solving Regularized Linear Programs Using Barrier Methods and KKT Systems*, Report SOL 96-4, Department of EESOR, Stanford University, Stanford, CA, 1996.
- [43] V. SIMONCINI AND D. B. SZYLD, *Flexible inner-outer Krylov subspace methods*, SIAM J. Numer. Anal., 40 (2003), pp. 2219–2239.
- [44] T. STEIHAUG, *The conjugate gradient method and trust regions in large scale optimization*, SIAM J. Numer. Anal., 20 (1983), pp. 626–637.
- [45] PH. L. TOINT, *Towards an efficient sparsity exploiting Newton method for minimization*, in Sparse Matrices and Their Uses, I. S. Duff, ed., Academic Press, London, New York, 1981, pp. 57–88.
- [46] R. J. VANDERBEI AND T. J. CARPENTER, *Symmetric indefinite systems for interior point methods*, Math. Program., 58 (1993), pp. 1–32.
- [47] R. J. VANDERBEI AND D. F. SHANNO, *An interior-point algorithm for nonconvex nonlinear programming*, Comput. Optim. Appl., 13 (1999), pp. 231–252.
- [48] S. J. WRIGHT AND D. ORBAN, *Properties of the log-barrier function on degenerate nonlinear programs*, Math. Oper. Res., 27 (2002), pp. 585–613.
- [49] S. J. WRIGHT, *Primal-Dual Interior-Point Methods*, SIAM, Philadelphia, PA, 1997.