SIAM J. OPTIM.
Vol. 24, No. 1, pp. 49–83

BOUNDS ON EIGENVALUES OF MATRICES ARISING FROM INTERIOR-POINT METHODS*

CHEN GREIF†, ERIN MOULDING‡, AND DOMINIQUE ORBAN§

Abstract. Interior-point methods feature prominently among numerical methods for inequality-constrained optimization problems, and involve the need to solve a sequence of linear systems that typically become increasingly ill-conditioned with the iterations. To solve these systems, whose original form has a nonsymmetric 3×3 block structure, it is common practice to perform block elimination and either solve the resulting reduced saddle-point system, or further reduce the system to the Schur complement and apply a symmetric positive definite solver. In this paper we use energy estimates to obtain bounds on the eigenvalues of the matrices, and conclude that the original unreduced matrix has more favorable eigenvalue bounds than the alternative reduced versions. Our analysis includes regularized variants of those matrices that do not require typical regularity assumptions.

Key words. primal-dual interior-point methods, convex quadratic optimization, indefinite linear systems, eigenvalues, condition number, inertia, eigenvalue bounds, regularization

AMS subject classifications. 15A06, 15A12, 15A18, 15B57, 90C06, 90C20, 65F22

DOI. 10.1137/120890600

1. Introduction.

Given a symmetric and positive semidefinite Hessian matrix $H \in \mathbb{R}^{n\times n}$, vectors $c \in \mathbb{R}^n$ and $b \in \mathbb{R}^m$, and a Jacobian matrix $J \in \mathbb{R}^{m\times n}$, where $m \le n$, consider the primal-dual pair of quadratic programs (QP) in standard form

$$\text{(1.1a)} \qquad \underset{x}{\text{minimize}}\ c^Tx + \tfrac{1}{2}x^THx \quad \text{subject to} \quad Jx = b, \quad x \ge 0,$$

$$\text{(1.1b)} \qquad \underset{x,y,z}{\text{maximize}}\ b^Ty - \tfrac{1}{2}x^THx \quad \text{subject to } J^Ty + z - Hx = c, \quad z \ge 0,$$

where inequalities are understood elementwise, and y and z are the vectors of Lagrange multipliers associated with the equality and nonnegativity constraints of (1.1a), respectively. The case $H = 0$ corresponds to the linear programming problem in standard form. Numerical methods for solving (1.1) include the class of widely successful primal-dual interior-point methods. Their distinctive feature is that they approximately follow a smooth path lying inside the primal-dual feasible set all the way to an optimal solution.

This paper focuses on the linear systems that form the core of the iterations of primal-dual interior-point methods. Specifically, the matrices associated with those linear systems have a special block form, and techniques that rely on partial elimination of the unknowns are fairly popular, as the underlying sparsity pattern naturally lends itself to such reductions. However, we claim that in terms of eigenvalues and conditioning, it may be beneficial to avoid performing such elimination steps before

*Received by the editors September 7, 2012; accepted for publication (in revised form) September 25, 2013; published electronically January 9, 2014.
http://www.siam.org/journals/siopt/24-1/89060.html

†Department of Computer Science, The University of British Columbia, Vancouver, BC, Canada (greif@cs.ubc.ca). This author's research was partially supported by an NSERC Discovery Grant.

‡Department of Mathematics, The University of British Columbia, Vancouver, BC, Canada (moulding@math.ubc.ca). This author's research was partially supported by an NSERC CGS M Scholarship.

§GERAD and Department of Mathematics and Industrial Engineering, École Polytechnique, Montréal, QC, Canada (dominique.orban@gerad.ca). This author's research was partially supported by an NSERC Discovery Grant.

applying a linear solver. Some of the linear systems with which we are concerned are nonsymmetric and diagonalizable, and eigenvalues describe the convergence behavior of certain iterative methods—see for example (Saad, 2003, Section 6.11) or (Nachtigal, Reddy, and Trefethen, 1992). To make our point, we use energy estimates in the spirit of Rusten and Winther (1992), and obtain upper and lower bounds on the eigenvalues of the various matrices that we consider. We also consider *regularized* variants of those matrices that arise when an interior-point method is applied to a modified optimization problem that includes regularization terms.

Our primary goal is to provide a theoretical foundation to the study of spectral properties of the matrices involved. When the problem is very large, the spectral structure of the matrices plays a central role in the performance of the interior-point solver, especially when the underlying linear systems are solved iteratively. We stress however that we do not perform analysis for iterative solvers or offer specialized preconditioning approaches in the present paper. We also note that the (standard) scaling approach that we use involves an ill-conditioned diagonal matrix, and its influence on the performance of iterative solvers warrants further investigation.

In section 2 we provide a short overview of primal-dual interior-point methods and present the linear systems that arise throughout the iterations. This section includes a regularized formulation of the optimization problem, which we will extensively analyze throughout the paper. In section 3 we specify conditions for nonsingularity and study the inertia of the matrices involved. In section 4 we analyze the regularized optimization problem introduced in section 2, provide bounds on the eigenvalues of the linear systems, and argue that in most cases, those bounds are tighter than existing related bounds. In section 5 we provide results for the original (unregularized) problem (1.1), which can be obtained as a special case of the analysis of the regularized problem. In section 6 we provide numerical validation of our analytical claims and in section 7 we cover several alternative system formulations. Concluding remarks appear in section 8.

2. Background and preliminaries. In this section we provide a brief overview of primal-dual interior-point methods and the linear systems that arise throughout the iterations. Our purpose is to set the stage for the analysis of the subsequent sections.

For an index set $\mathcal{N} \subseteq \{1, \ldots, n\}$ and a vector $v \in \mathbb{R}^n$ we denote by $v_{\mathcal{N}}$ the subvector of v indexed by $\mathcal{N}$. Similarly, if A is a matrix with n columns, $A_{\mathcal{N}}$ is the submatrix of the columns of A corresponding to indices in $\mathcal{N}$. If A is square, $A_{\mathcal{N}\mathcal{N}}$ represents the matrix with both rows and columns corresponding to indices in $\mathcal{N}$.

Throughout the paper we separate vector components by commas. To avoid ambiguities, inner products are denoted by a transpose operation. Thus, (x, y) is a vector whose components are x and y (each of which may be a vector) whereas $x^T y$ is the inner product of the vectors x and y.

2.1. Primal-dual interior-point methods. If x is feasible for (1.1a), we let

$$\mathcal{A}(x) := \{i = 1, \ldots, n \mid x_i = 0\} \quad \text{and} \quad \mathcal{I}(x) := \{1, \ldots, n\} \setminus \mathcal{A}(x)$$

be the index sets of active and inactive bounds, respectively. For simplicity and when there is no ambiguity, we write $\mathcal{A}$ and $\mathcal{I}$ instead of $\mathcal{A}(x)$ and $\mathcal{I}(x)$. All solutions (x, y, z) of (1.1) must satisfy the complementarity condition

$$x_i z_i = 0 \quad \text{for all} \quad i = 1, \ldots, n,$$

which may also be written $z_{\mathcal{I}} = 0$. A solution (x, y, z) of (1.1) is *strictly complementary* if $z_i > 0$ for all $i \in \mathcal{A}$, which may also be written $z_{\mathcal{A}} > 0$.

It is generally assumed that Slater's constraint qualification condition holds, i.e., that there exists a primal-dual triple (x, y, z) such that $Jx = b$, $J^T y + z - Hx = c$, and $(x, z) > 0$. Primal-dual interior-point methods generate a sequence of iterates (x_k, y_k, z_k) that remain strictly feasible with respect to the bound constraints, i.e., $(x_k, z_k) > 0$, but not necessarily to the equality constraints, with the intention of satisfying, in the limit, the common necessary and sufficient first-order optimality conditions of (1.1a) and (1.1b). Those iterates (approximately) satisfy the perturbed optimality conditions

$$\begin{bmatrix} c + Hx - J^T y - z \\ Jx - b \\ \tau e - XZe \end{bmatrix} = 0, \qquad (x, z) > 0. \tag{2.1}$$

Here the triple (x, y, z) represents a generic current iterate. We drop the subscript k for brevity and we use the standard notation X and Z to denote diagonal matrices whose diagonal elements are the components of the vectors x and z, respectively, and e to represent the vector of ones of appropriate size. The system depends on $\tau > 0$, the *barrier parameter*, which governs the progress of the interior-point method and converges to zero. This parameter is typically set as the product $\tau := \sigma\mu$, where

$$\mu := \frac{x^T z}{n}$$

is the *duality measure* and $\sigma \in [0, 1]$ is the *centering parameter* used to achieve a desirable reduction in the duality measure at each iteration. Tied to these parameters is the notion of *central path*, which is the set of triples $(x, y, z) = (x_\tau, y_\tau, z_\tau)$ of exact solutions to (2.1) associated with $\tau > 0$. Practical methods typically seek a compromise between improving centrality by taking a relatively large value of σ, which allows for taking a longer step in the next iteration, and reducing the duality measure μ by taking a small value of σ, which results in a direction closer to the pure Newton direction, often called the affine-scaling direction.

The matrix associated with the interior-point iterations is the Jacobian of the system of equations (2.1). It is of size $(2n + m) \times (2n + m)$ and can be most naturally written in a block 3×3 form. The linear system is given by

$$\begin{bmatrix} H & J^T & -I \\ J & 0 & 0 \\ -Z & 0 & -X \end{bmatrix} \begin{bmatrix} \Delta x \\ -\Delta y \\ \Delta z \end{bmatrix} = \begin{bmatrix} -c - Hx + J^T y + z \\ b - Jx \\ XZe - \tau e \end{bmatrix}. \tag{2.2}$$

Most blocks of the matrix stay fixed throughout the interior-point iterations. The only ones that change, X and Z, are diagonal with strictly positive diagonal elements during the iteration, although in the limit some diagonal elements typically vanish.

We note that there are several ways to arrange the system (2.2) in terms of signs and ordering of unknowns; we choose this formulation in anticipation of symmetrizing the matrix in the form of a suggestion of Saunders—see (Forsgren, 2002). The barrier parameter appears explicitly only in the right-hand side, but it also influences the matrix itself since the iterates x and z, if they converge, do so at an asymptotic rate that is a function of the duality measure μ. The solution of (2.2) for $(\Delta x, \Delta y, \Delta z)$ defines the next iterate

$$(x^+, y^+, z^+) = (x, y, z) + \alpha(\Delta x, \Delta y, \Delta z),$$

where $\alpha \in (0, 1]$ is a step length chosen to ensure that $(x^+, z^+) > 0$ and possibly other conditions.

2.2. Block elimination approaches. Given the block structure of the matrix of (2.2), a few possibilities for solving the system naturally arise due to the special sparsity structure, particularly the diagonality and positive definiteness of X and Z. An obvious (though not common) approach is that of directly solving the linear system (2.2). The matrix is mildly nonsymmetric but easily symmetrizable (see section 3), and so it is possible to apply symmetric solvers. We note, however, that the symmetrization is done by a diagonal matrix that becomes increasingly ill-conditioned as iterations progress, and hence convergence and numerical accuracy may be affected.

A second approach is that of exploiting the nonsingularity and diagonality of X to perform one step of block Gaussian elimination and obtain

$$\begin{bmatrix} H + X^{-1}Z & J^T \\ J & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} -c - Hx + J^T y + \tau X^{-1} e \\ b - Jx \end{bmatrix}. \tag{2.3}$$

The matrix of (2.3) is a typical symmetric indefinite saddle-point matrix and it has size $(n + m) \times (n + m)$. Significant progress has been made on numerical solution methods for saddle-point systems in the past few decades (Benzi, Golub, and Liesen, 2005), but here a specific challenge is that $X^{-1}Z$ may cause ill-conditioning as X and Z have diagonal entries iterating toward zero. The stability of the symmetric indefinite factorization of the matrix of (2.3) has been studied in Forsgren, Gill, and Shinnerl (1996).

A third approach is to take an additional step of block Gaussian elimination before applying a linear solver. This amounts to forming the Schur complement equations

$$J(H + X^{-1}Z)^{-1}J^T \Delta y = b - Jx - J(H + X^{-1}Z)^{-1}(-c - Hx + J^T y + \tau X^{-1} e). \tag{2.4}$$

The matrix associated with the linear system (2.4) is positive definite provided J has full row rank. This approach is popular with practitioners, since symmetric positive definite solvers are often preferred over indefinite solvers, and the Schur complement equations are smaller, of size $m \times m$. However, in cases other than linear programming, forming system (2.4) comes at the potentially costly price of having to first invert or factor $H + X^{-1}Z$ and having to deal with a significant loss of sparsity.

Since the matrices associated with the linear systems (2.2), (2.3), and (2.4) feature prominently in the discussion that ensues throughout the paper, we denote them as follows, where the subscript stands for the "natural" block size of the matrix:

$$K_3 := \begin{bmatrix} H & J^T & -I \\ J & 0 & 0 \\ -Z & 0 & -X \end{bmatrix}, \tag{2.5a}$$

$$K_2 := \begin{bmatrix} H + X^{-1}Z & J^T \\ J & 0 \end{bmatrix}, \tag{2.5b}$$

$$K_1 := J(H + X^{-1}Z)^{-1}J^T. \tag{2.5c}$$

2.3. Regularization. Numerical difficulties may arise if J does not have full row rank, if strict complementarity is not satisfied in the limit—see section 2.1—or when the linear independence qualification condition is not satisfied—see Definition 2.2 in the next section and a relevant discussion in section 3.

One way to alleviate some of those difficulties is by introducing regularization terms (Saunders, 1996). This may be done in a variety of ways and we focus on a

two-parameter regularization approach, aimed at taking eigenvalues of the Hessian and singular values of the Jacobian away from zero. See, e.g., Gondzio (2012) for a similar approach. We introduce parameters $\rho > 0$ and $\delta > 0$, and consider an *exact regularization* approach proposed by Friedlander and Orban (2012) for the primal QP problem:

$$\begin{aligned} \underset{x,r}{\text{minimize}} \quad & c^T x + \tfrac{1}{2} x^T H x + \tfrac{1}{2}\rho\|x - x_k\|^2 + \tfrac{1}{2}\delta\|r + y_k\|^2 \\ \text{subject to} \quad & Jx + \delta r = b,\ x \geq 0. \end{aligned} \tag{2.6}$$

Here x_k and y_k are current primal and dual approximations, respectively. The corresponding dual problem is given as

$$\begin{aligned} \underset{x,y,z,s}{\text{maximize}} \quad & b^T y - \tfrac{1}{2} x^T H x - \tfrac{1}{2}\delta\|y - y_k\|^2 - \tfrac{1}{2}\rho\|s + x_k\|^2 \\ \text{subject to} \quad & -Hx + J^T y + z - \rho s = c,\ z \geq 0. \end{aligned} \tag{2.7}$$

Note that setting $\delta = \rho = 0$ recovers the original primal-dual pair (1.1a)–(1.1b).

Friedlander and Orban (2012) propose an interior-point method for (2.6)–(2.7) that converges under standard conditions with either fixed or decreasing values of the regularization parameters, without assumptions on the rank of J.

For the regularized approach (2.6)–(2.7) the associated linear systems involve a modified version of the 3×3 block system (2.2), with the right-hand side unchanged and the matrix given as follows:

$$K_{3,\text{reg}} := \begin{bmatrix} H + \rho I & J^T & -I \\ J & -\delta I & 0 \\ -Z & 0 & -X \end{bmatrix}. \tag{2.8a}$$

Upon reduction, the 2×2 matrix reads

$$K_{2,\text{reg}} := \begin{bmatrix} H + X^{-1}Z + \rho I & J^T \\ J & -\delta I \end{bmatrix}. \tag{2.8b}$$

Finally, the Schur complement equations now have the matrix

$$K_{1,\text{reg}} := J(H + X^{-1}Z + \rho I)^{-1} J^T + \delta I. \tag{2.8c}$$

Our results essentially consider two distinct situations to analyze the properties of K_1, K_2 and K_3 and, respectively, of $K_{1,\text{reg}}$, $K_{2,\text{reg}}$, and $K_{3,\text{reg}}$. The first concerns values of the matrix throughout the iterations while the second concerns the value in the limit, at a point satisfying complementarity, such as a solution of (1.1).

2.4. Notation and further definitions. We denote the eigenvalues of the Hessian H by $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq 0$, and the singular values of J by $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_m \geq 0$. The spectral condition number of a matrix B, defined as $\sigma_{\max}(B)/\sigma_{\min}(B)$, will be denoted by $\kappa(B)$, where $\sigma_{\max}$ and $\sigma_{\min}$ denote the largest and smallest singular values, respectively. We use the notation γ and θ to denote a generic eigenvalue of (2.8b) and (2.8a), respectively.

For two related positive scalar sequences $\{\alpha_k\}$ and $\{\beta_k\}$, we write $\alpha_k = O(\beta_k)$ if there exists a constant $\varphi > 0$ such that $\alpha_k \leq \varphi\beta_k$ for all sufficiently large k. Equivalently, we write $\beta_k = \Omega(\alpha_k)$. We write $\alpha_k = \Theta(\beta_k)$ if $\alpha_k = O(\beta_k)$ and $\beta_k = O(\alpha_k)$. If $\alpha_k = O(\beta_k)$ and $\alpha_k = \Omega(\gamma_k)$ for some third positive sequence $\{\gamma_k\}$, we use the shorthand notation $\gamma_k \lesssim \alpha_k \lesssim \beta_k$.

Throughout our analysis, the following two definitions will be useful.

DEFINITION 2.1 (inertia). *For a given symmetric matrix M, the inertia of M is the triple (n^+, n^-, n^0), where n^+, n^-, and n^0 are the numbers of positive, negative, and zero eigenvalues of M, respectively.*

We note that the definition of inertia extends to nonsymmetric matrices with real eigenvalues, but our focus is on symmetric matrices. Given a symmetric matrix M, if C is a real nonsingular matrix and $N = CMC^T$, Sylvester's law of inertia asserts that N and M have the same inertia.

The following definition states a standard qualification condition required to ensure certain nonsingularity properties.

DEFINITION 2.2 (LICQ). *The linear independence constraint qualification condition is satisfied at x, feasible for* (1.1a)*, if $\begin{bmatrix} J^T & -I_{\mathcal{A}(x)} \end{bmatrix}$ has full column rank.*

Note that the LICQ imposes an upper bound on the size of the active set: $|\mathcal{A}(x)| \leq n - m$. If the LICQ and strict complementarity are satisfied at a solution, we say that this solution is nondegenerate. As we shall see, these conditions guarantee that the matrices of interest are nonsingular. These are common assumptions in optimization.

Throughout the paper, we illustrate our bounds on the following generic, but typical, example situation.

Example 2.1. We consider a generic interior-point method guaranteeing the following asymptotic estimates:

$$x_i = \Theta(\mu) \quad (i \in \mathcal{A}), \qquad x_i = \Theta(1) \quad (i \in \mathcal{I}), \tag{2.9a}$$

$$z_i = \Theta(\mu) \quad (i \in \mathcal{I}), \qquad z_i = \Theta(1) \quad (i \in \mathcal{A}). \tag{2.9b}$$

We assume that $\mathcal{A} \neq \varnothing$ and $\mathcal{I} \neq \varnothing$.

Most problems are such that $\mathcal{A} \neq \varnothing$ and $\mathcal{I} \neq \varnothing$ and most interior-point methods applied to a nondegenerate problem match the situation of Example 2.1. In particular, when (x, y, z) are exact solutions of (2.1), we have $X^{-1}Z = \mu X^{-2} = \Theta(1/\mu)$, but this estimate also holds sufficiently close to the central path. Indeed most interior-point algorithms for convex quadratic programming confine the iterates to a neighborhood of the central path defined, among other conditions, by the requirement that $\gamma_1\mu \leq x_i z_i \leq \gamma_2\mu$ for all $i = 1, \ldots, n$, for some positive constants γ_1 and γ_2. That the relations in (2.9) hold under strict complementarity is then a simple consequence; see, e.g., (Wright, 1997).

In our implementation we use the predictor-corrector scheme due to Mehrotra (1992), which is based on first taking the pure Newton direction, i.e., with $\sigma = 0$, and then following a step aiming toward the central path as a correction for the linearization error in XZ. The algorithm thus solves two linear systems with the same matrix but with different right-hand sides. Although this algorithm does not confine the iterates to a neighborhood of the sort mentioned above, we will assume that (2.9) holds.

We formalize now a few working assumptions related to convexity, positivity, and complementarity, and refer to them throughout the paper. We note that the last two assumptions are mutually exclusive.

Assumption 2.1 (convexity). The Hessian H is symmetric and positive semidefinite.
Assumption 2.2 (positivity). The triple (x, y, z) satisfies $(x, z) > 0$.
Assumption 2.3 (complementarity). The triple (x, y, z) satisfies $(x, z) \geq 0$ and $x_i z_i = 0$ for all $i = 1, \ldots, n$.

2.5. Related work. While the algorithms of modern interior-point solvers are mostly settled, the choice of linear system formulation differs across software packages. Many modern solvers reduce to the Schur complement equations form, e.g., PCx for linear programming (Czyzyk et al., 1999). Others reduce to the saddle-point form, e.g., OOQP for quadratic programming (Gertz and Wright, 2003) and IPOPT and KNITRO for general nonlinear programming (Byrd, Hribar, and Nocedal, 1999; Byrd, Gilbert, and Nocedal, 2000; Wächter and Biegler, 2006). Another example is HOPDM for linear programming and convex QP, which automatically chooses either the Schur complement equations or saddle-point form (Altman and Gondzio, 1999). We are not aware of existing solvers that solve the unreduced system (2.2) for any of these problems.

Gondzio (2012) provides analysis of properties of the Schur complement equations form (2.4). The saddle-point formulation (2.3) has properties that directly follow from the general properties of such matrices—see (Rusten and Winther, 1992; Benzi, Golub, and Liesen, 2005; Silvester and Wathen, 1994; Gould and Simoncini, 2010; Axelsson and Neytcheva, 2006; Bai, Ng, and Wang, 2009) for some relevant general results, and (Friedlander and Orban, 2012) for results specialized to optimization. The ill-conditioning of some reduced matrices is well known (Fiacco and McCormick, 1990; Wright, 1992, 2005; Forsgren, 2002; Forsgren, Gill, and Wright, 2002), but it has been referred to, with some assumptions on solution methods, as "benign" (Wright, 2005; Forsgren, 2002), "usually harmless" (Forsgren, Gill, and Wright, 2002), and "highly structured" (Forsgren, Gill, and Wright, 2002). The matrices for classical barrier methods, corresponding to the choice $Z = \tau X^{-1}$, are also ill-conditioned (Wright, 1994; Forsgren, Gill, and Wright, 2002).

There exist relatively few results on the unreduced 3×3 formulation. Korzak (1999) covers some spectral properties of various formulations for the special case of linear programming. Armand and Benoist (2011) prove uniform boundedness of the inverse under several assumptions, intended to be used in further theoretical analysis. A private communication of Saunders is cited by Forsgren (2002) who notes the symmetrizability and potential appeal of the 3×3 system, equivalent to the symmetrized matrix used in this paper. Forsgren (2002) and Forsgren, Gill, and Wright (2002) note that the matrix of this system remains well-conditioned though ill-conditioning remains an issue when forming the right-hand side and unscaling the variables, due to multiplication by a diagonal matrix with large elements; these papers mention also a different (ill-conditioned) symmetric formulation. Finally, it is possible to represent the 3×3 block matrix as a 2×2 saddle-point matrix by splitting appropriately, and use known results to obtain eigenvalue bounds—see, for example, Bai (2013). This type of splitting is examined in section 4.3 and is used to establish that our results are generally tighter than bounds resulting from the application of such known results.

3. Nonsingularity and inertia. In this section we specify conditions for nonsingularity of the various matrices, and determine the inertia during the iterations and at the limit. We rely on techniques introduced by Gould (1985) and Forsgren (2002) to prove our results, introduce a series of propositions, and conclude the section

with two results for the 3×3 block case. See also Moulding (2012) for an additional discussion and a few alternative proofs.

3.1. Preliminary results. We begin with two fundamental technical results.

LEMMA 3.1 (Forsgren, 2002, Proposition 2). *Let $A = A^T \in \mathbb{R}^{q\times q}$, $B \in \mathbb{R}^{t\times q}$, $C = C^T \in \mathbb{R}^{t\times t}$ positive semidefinite,*

$$K := \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix},$$

and $r := \text{Rank}(\begin{bmatrix} B & -C \end{bmatrix})$. Let the columns of U form a basis for $\text{Null}(C)$, the columns of N form a basis for $\text{Null}(U^T B)$ and p be the dimension of $\text{Null}(C)$. Finally, let $C^\dagger$ denote the pseudoinverse of C. Then

$$\text{In}(K) = \text{In}\left(N^T(A + B^T C^\dagger B)N\right) + (p - t + r,\, r,\, t - r).$$

In addition, $\text{Rank}(U^T B) = p - t + r$.

When $C = 0$, Lemma 3.1 reduces to Lemma 3.2.

LEMMA 3.2 (Gould, 1985, Lemma 3.4). *Let $A = A^T \in \mathbb{R}^{q\times q}$, $B \in \mathbb{R}^{t\times q}$, and*

$$K := \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix}.$$

Let $r := \text{Rank}(B)$ and the columns of N form a basis for $\text{Null}(B)$. Then

$$\text{In}(K) = \text{In}(N^T A N) + (r,\, r,\, t - r).$$

Lemma 3.1 and Lemma 3.2 may be employed to analyze the inertia—and therefore nonsingularity—of K_2, K_3, and their regularized variants. This is the subject of the rest of this section.

3.2. Inertia and nonsingularity of $\mathbf{K_2}$ and $\mathbf{K_{2,reg}}$. We begin our investigation with the 2×2 block systems. Strictly speaking, our first result does not require Assumption 2.1 (convexity) to hold, although of course, the positive semidefiniteness assumption on $H + \rho I$ may be removed if the convexity assumption holds. In other words, the following proposition holds also if H is indefinite.

PROPOSITION 3.3. *If Assumption 2.2 (positivity) holds, $\rho \geq 0$ is such that $H + \rho I$ is positive semidefinite, and $\delta > 0$, the inertia of $K_{2,\text{reg}}$ is $(n, m, 0)$.*

Proof. This follows from Lemma 3.1 with $q = n$, $A = H + X^{-1}Z + \rho I$, $t = m$, $B = J$, $C = \delta I$, $r = m$, the observations that $H + X^{-1}Z + \rho I$ and δI are positive definite, that $p = 0$, and that $C^\dagger = \delta^{-1} I$. In this case, U is vacuous and the columns of N form a basis for $\text{Null}(J)$. ∎

We may formulate a result concerning the nonsingularity of $K_{2,\text{reg}}$ as a direct consequence of Proposition 3.3.

COROLLARY 3.4. *If Assumption 2.2 (positivity) holds, $\rho \geq 0$ is such that $H + \rho I$ is positive semidefinite, and $\delta > 0$, $K_{2,\text{reg}}$ is nonsingular.*

Note that Corollary 3.4 is a special case of (Benzi, Golub, and Liesen, 2005, Theorem 3.1). When the system is not regularized, the results are identical but require Assumption 2.1 (convexity) to hold and J to have full row rank.

PROPOSITION 3.5. *If J has full row rank and Assumptions* 2.1 *(convexity) and* 2.2 *(positivity) are satisfied, the inertia of K_2 is $(n, m, 0)$.*

Proof. This follows directly from Lemma 3.2 with $q = n$, $A = H + X^{-1}Z$, $t = m$, $B = J$, $r = m$, and the observation that $H + X^{-1}Z$ is positive definite. ∎

The following nonsingularity result is a special case of (Benzi, Golub, and Liesen, 2005, Theorem 3.2). Part of the result may be viewed as a direct consequence of Proposition 3.5.

COROLLARY 3.6. *Suppose Assumptions* 2.1 *(convexity) and* 2.2 *(positivity) are satisfied. The matrix K_2 is nonsingular if and only if J has full row rank.*

3.3. Inertia and nonsingularity of $\mathbf{K_3}$ and $\mathbf{K_{3,reg}}$. It is useful to consider a symmetrized version of K_3, making it possible to work with real arithmetic. Indeed, it is easy to show that K_3 is symmetrizable and has real eigenvalues; see (Forsgren, 2002). Consider the diagonal matrix

$$D = \begin{bmatrix} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & Z^{\frac{1}{2}} \end{bmatrix}. \tag{3.1}$$

Using the similarity transformation associated with D, we obtain the symmetric matrix

$$\hat{K}_{3,\mathrm{reg}} := D^{-1} K_{3,\mathrm{reg}} D = \begin{bmatrix} H + \rho I & J^T & -Z^{\frac{1}{2}} \\ J & -\delta I & 0 \\ -Z^{\frac{1}{2}} & 0 & -X \end{bmatrix}. \tag{3.2}$$

When $\rho = \delta = 0$, the matrix of (3.2) is denoted $\hat{K}_3$.

We begin with results on the inertia of $\hat{K}_{3,\mathrm{reg}}$ both during the iterations and in the limit. As before, Assumption 2.1 (convexity) is not required to hold. Note that if it does hold and $\rho > 0$, the assumption on the intersection of nullspaces in Lemma 3.7 automatically holds because $H + \rho I$ is positive definite.

LEMMA 3.7. *Let $\rho \geq 0$, $\delta > 0$ and assume that $H + \rho I$ is positive semidefinite. Suppose that either*

1. *Assumption* 2.2 *(positivity) holds, or*
2. *Assumption* 2.3 *(complementarity) holds at (x, y, z), where strict complementarity is satisfied and* $\mathrm{Null}(H + \rho I) \cap \mathrm{Null}(J) \cap \mathrm{Null}(Z) = \{0\}$.

Then $\mathrm{In}(\hat{K}_{3,\mathrm{reg}}) = (n,\, n + m,\, 0)$.

Proof. Suppose first that Assumption 2.2 (positivity) holds. The block decomposition

$$\hat{K}_{3,\mathrm{reg}} = \begin{bmatrix} I & & Z^{\frac{1}{2}}X^{-1} \\ & I & \\ & & I \end{bmatrix} \begin{bmatrix} H + X^{-1}Z + \rho I & J^T & \\ J & -\delta I & \\ & & -X \end{bmatrix} \begin{bmatrix} I & & \\ & I & \\ X^{-1}Z^{\frac{1}{2}} & & I \end{bmatrix} \tag{3.3}$$

and Sylvester's law of inertia show that $\mathrm{In}(\hat{K}_{3,\mathrm{reg}}) = \mathrm{In}(K_{2,\mathrm{reg}}) + \mathrm{In}(-X) = (n,\, n + m,\, 0)$.

Suppose now that Assumption 2.3 (complementarity) is satisfied and furthermore, strict complementarity is satisfied at (x, y, z). Then $x_{\mathcal{A}} = 0$, $x_{\mathcal{I}} > 0$, $z_{\mathcal{A}} > 0$, and

$z_{\mathcal{I}} = 0$. Using the partitioning induced by $\mathcal{A}$ and $\mathcal{I}$, we may write

$$\hat{K}_{3,\text{reg}} = \begin{bmatrix} H^{\rho}_{\mathcal{A}\mathcal{A}} & H^{\rho}_{\mathcal{A}\mathcal{I}}{}^{T} & J_{\mathcal{A}}^{T} & -Z_{\mathcal{A}}^{\frac{1}{2}} & \\ H^{\rho}_{\mathcal{A}\mathcal{I}} & H^{\rho}_{\mathcal{I}\mathcal{I}} & J_{\mathcal{I}}^{T} & & \\ J_{\mathcal{A}} & J_{\mathcal{I}} & -\delta I & & \\ -Z_{\mathcal{A}}^{\frac{1}{2}} & & & & \\ & & & & -X_{\mathcal{I}} \end{bmatrix}, \tag{3.4}$$

where $H^{\rho} := H + \rho I$. In order to apply Lemma 3.1, we let $n_{\mathcal{A}} := |\mathcal{A}|$, $n_{\mathcal{I}} := |\mathcal{I}|$, and define

$$B := \begin{bmatrix} J_{\mathcal{A}} & J_{\mathcal{I}} \\ -Z_{\mathcal{A}}^{\frac{1}{2}} & 0 \\ 0 & 0 \end{bmatrix}, \quad C := \begin{bmatrix} \delta I & & \\ & 0 & \\ & & X_{\mathcal{I}} \end{bmatrix}, \quad U := \begin{bmatrix} 0_{m\times n_{\mathcal{A}}} \\ I_{n_{\mathcal{A}}\times n_{\mathcal{A}}} \\ 0_{n_{\mathcal{I}}\times n_{\mathcal{A}}} \end{bmatrix}, \quad N := \begin{bmatrix} 0_{n_{\mathcal{A}}\times n_{\mathcal{I}}} \\ I_{n_{\mathcal{I}}\times n_{\mathcal{I}}} \end{bmatrix},$$

where we indicate the block dimensions for clarity. Since $\delta > 0$ it is clear that $\text{Rank}(\begin{bmatrix} B & -C \end{bmatrix}) = n + m$. Note that the $p = n_{\mathcal{A}}$ columns of U form a basis for $\text{Null}(C)$ while those of N form a basis for $\text{Null}(U^T B)$. Because $C^{\dagger}$ is the block diagonal matrix $\text{diag}(\delta^{-1} I, 0, X_{\mathcal{I}}^{-1})$, we have

$$N^T (H^{\rho} + B^T C^{\dagger} B) N = N^T (H^{\rho} + \tfrac{1}{\delta} J^T J) N = H^{\rho}_{\mathcal{I}\mathcal{I}} + \tfrac{1}{\delta} J_{\mathcal{I}}^T J_{\mathcal{I}}.$$

This last matrix is positive definite because $\text{Null}(H + \rho I) \cap \text{Null}(J) \cap \text{Null}(Z) = \{0\}$. Indeed, the columns of

$$N_Z := \begin{bmatrix} 0_{n_{\mathcal{A}}\times n_{\mathcal{I}}} \\ I_{n_{\mathcal{I}}\times n_{\mathcal{I}}} \end{bmatrix}$$

form a basis for $\text{Null}(Z)$, and $N_Z^T (H^{\rho} + \frac{1}{\delta} J^T J) N_Z = H^{\rho}_{\mathcal{I}\mathcal{I}} + \frac{1}{\delta} J_{\mathcal{I}}^T J_{\mathcal{I}}$. Lemma 3.1 then yields

$$\text{In}(\hat{K}_{3,\text{reg}}) = (n_{\mathcal{I}},\, 0,\, 0) + (n_{\mathcal{A}} - (n+m) + (n+m),\, n+m,\, 0) = (n,\, n+m,\, 0),$$

which completes the proof. ∎

We now give the inertia of $\hat{K}_3$ and show that it is the same as in the regularized case both during the iterations and in the limit.

LEMMA 3.8. *Suppose Assumption* 2.1 *(convexity) holds and either*

1. *Assumption* 2.2 *(positivity) is satisfied and J has full row rank, or*
2. *Assumption* 2.3 *(complementarity) holds at (x, y, z), where strict complementarity holds, $J_{\mathcal{I}}$ has full row rank, and* $\text{Null}(H) \cap \text{Null}(J) \cap \text{Null}(Z) = \{0\}$.

Then $\text{In}(\hat{K}_3) = (n,\, n+m,\, 0)$.

Proof. Under Assumption 2.2 (positivity), if J has full row rank, then (3.3) is still valid, so that $\text{In}(\hat{K}_3) = \text{In}(K_2) + \text{In}(-X) = (n,\, n+m,\, 0)$.

Under Assumption 2.3 (complementarity), we repeat the proof of Lemma 3.7 using this time

$$C := \begin{bmatrix} 0 & & \\ & 0 & \\ & & X_{\mathcal{I}} \end{bmatrix}, \quad U := \begin{bmatrix} I_{m\times m} & 0_{m\times n_{\mathcal{A}}} \\ 0_{n_{\mathcal{A}}\times m} & I_{n_{\mathcal{A}}\times n_{\mathcal{A}}} \\ 0_{n_{\mathcal{I}}\times m} & 0_{n_{\mathcal{I}}\times n_{\mathcal{A}}} \end{bmatrix}, \quad N := \begin{bmatrix} 0_{n_{\mathcal{A}}\times(n_{\mathcal{I}}-m)} \\ N_{\mathcal{I}} \end{bmatrix},$$

where the $(n_{\mathcal{I}} - m)$ columns of $N_{\mathcal{I}}$ form a basis for $\text{Null}(J_{\mathcal{I}})$. Because $J_{\mathcal{I}}$ has full row rank, we must have $n_{\mathcal{I}} \geq m$. Note that the columns of U form a basis for $\text{Null}(C)$ while those of N form a basis for $\text{Null}(U^T B)$. This time,

$$N^T(H + B^T C^\dagger B)N = N^T H N = N_{\mathcal{I}}^T H_{\mathcal{I}\mathcal{I}} N_{\mathcal{I}},$$

which is positive definite because $\text{Null}(H) \cap \text{Null}(J) \cap \text{Null}(Z) = \{0\}$. Indeed, the columns of

$$N_Z := \begin{bmatrix} 0_{n_{\mathcal{A}} \times n_{\mathcal{I}}} \\ I_{n_{\mathcal{I}} \times n_{\mathcal{I}}} \end{bmatrix}$$

form a basis for $\text{Null}(Z)$, and $N_Z^T H N_Z = H_{\mathcal{I}\mathcal{I}}$. Lemma 3.1 now yields

$$\text{In}(\hat{K}_3) = (n_{\mathcal{I}} - m,\, 0,\, 0) + (m + n_{\mathcal{A}} - (n+m) + (n+m),\, n+m,\, 0) = (n,\, n+m,\, 0),$$

which completes the proof. □

Lemma 3.8 also holds if H is indefinite yet positive definite on the nullspace of J. This corresponds more closely to a typical case in a neighborhood of an isolated minimizer, where H is positive definite over a subset of $\text{Null}(J)$ only—a cone defined by the gradients of the active bounds.

Since singularity occurs simultaneously for K_3 and $\hat{K}_3$, and their regularized counterparts, either of them can be considered, and we choose to work with the nonsymmetric versions. The results above give sufficient conditions for K_3 and $K_{3,\text{reg}}$ to be nonsingular. In the rest of this section, we give circumstances under which those conditions are also necessary. We begin with conditions for $K_{3,\text{reg}}$ to be nonsingular throughout the interior-point iteration. The next result also covers the unregularized setting $\delta = \rho = 0$.

PROPOSITION 3.9. *Let $\rho \geq 0$ and suppose Assumptions* 2.1 *(convexity) and* 2.2 *(positivity) hold. The matrix $K_{3,\text{reg}}$ is nonsingular if and only if either* (i) $\delta > 0$ *or* (ii) $\delta = 0$ *and J has full rank.*

Proof. Since X is nonsingular, $K_{3,\text{reg}}$ is nonsingular if and only if the matrix

$$\begin{bmatrix} H + \rho I + X^{-1}Z & J^T \\ J & -\delta I \end{bmatrix}$$

is nonsingular. Case (i), namely, $\delta > 0$, follows either by applying Lemma 3.1 or the fact that the above displayed matrix is symmetric quasidefinite (Vanderbei, 1995) since $H + \rho I + X^{-1}Z$ is symmetric positive definite. Case (ii), namely, $\delta = 0$, follows by Lemma 3.2; nonsingularity holds if and only if J has full rank. □

As mentioned earlier, inspection of the proof of Proposition 3.9 reveals that our initial assumption of positive semidefinite H may be weakened. Indeed, it is sufficient to assume that H is positive semidefinite on the nullspace of J only. In this case however, the duality relationship between (1.1a) and (1.1b) is no longer so simple. Nevertheless, such a restricted definiteness assumption is classic in nonconvex optimization—see, e.g., Gould (1985).

We now consider what happens to K_3 in the limit of the interior-point method.

PROPOSITION 3.10. *Suppose Assumptions* 2.1 *(convexity) and* 2.3 *(complementarity) hold at (x, y, z). Then K_3 is nonsingular if and only if the solution (x, y, z) is strictly complementary,* $\text{Null}(H) \cap \text{Null}(J) \cap \text{Null}(Z) = \{0\}$, *and the LICQ is satisfied.*

Proof. If (x, y, z) is not strictly complementary, there is a zero row in the third block row of (2.2) and K_3 is singular. Therefore, strict complementarity is necessary.

Consider the system

$$\begin{bmatrix} H & J^T & -I \\ J & 0 & 0 \\ -Z & 0 & -X \end{bmatrix} \begin{bmatrix} u \\ v \\ w \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \tag{3.5}$$

If $\text{Null}(H) \cap \text{Null}(J) \cap \text{Null}(Z) \neq \{0\}$, take $0 \neq u \in \text{Null}(H) \cap \text{Null}(J) \cap \text{Null}(Z)$, $v = 0$, and $w = 0$. Since $Hu = Ju = -Zu = 0$, it follows from (3.5) that (u, v, w) is a nontrivial null vector of K_3. Thus this condition is necessary.

Now, assume strict complementarity and $\text{Null}(H) \cap \text{Null}(J) \cap \text{Null}(Z) = \{0\}$, and suppose (u, v, w) is in the nullspace of K_3. Since $z_{\mathcal{I}} = 0$ at the solution (see section 2.1) by complementarity, we have $Z = \text{diag}\,(z_{\mathcal{A}}, z_{\mathcal{I}}) = \text{diag}\,(z_{\mathcal{A}}, 0)$, with $z_{\mathcal{A}} > 0$, so that $\text{Null}(Z) = \text{span}\{e_i | i \in \mathcal{I}\}$. The third block row of (3.5) and strict complementarity necessarily yield $u_{\mathcal{A}} = 0$ and $w_{\mathcal{I}} = 0$, and so $u = (0, u_{\mathcal{I}})$, which implies that u lies entirely in the nullspace of Z. Therefore, $u^T w = 0$. Taking the inner product of the first block row of (3.5) with u and substituting $Ju = 0$ from the second block row gives $u^T H u = 0$. We must thus have $u \in \text{Null}(H) \cap \text{Null}(J) \cap \text{Null}(Z)$, which implies that $u = 0$. Eliminating u and $w_{\mathcal{I}}$ from (3.5), we have

$$\begin{bmatrix} J^T & -I_{\mathcal{A}} \end{bmatrix} \begin{bmatrix} v \\ w_{\mathcal{A}} \end{bmatrix} = 0,$$

which has only the trivial solution if and only if the LICQ holds. ∎

Finally, we consider what happens to $K_{3,\text{reg}}$ in the limit of the interior-point iteration. If (x, y, z) is not strictly complementary, there is a zero row in the third block row of (2.8a) and $K_{3,\text{reg}}$ is singular. Thus strict complementarity is necessary for nonsingularity in each case. The proposition below includes the unregularized case, $\delta = \rho = 0$. The proof is similar to that of Proposition 3.10.

PROPOSITION 3.11. *Suppose Assumptions* 2.1 *(convexity) and* 2.3 *(complementarity) hold at* (x, y, z)*. Necessary and sufficient conditions for the matrix* $K_{3,\text{reg}}$ *to be nonsingular are that* (x, y, z) *be strictly complementary, and*

- $\text{Null}(H) \cap \text{Null}(J) \cap \text{Null}(Z) = \{0\}$ *if* $\rho = 0$*, and*
- *the LICQ be satisfied if* $\delta = 0$*.*

4. Eigenvalue bounds for the regularized systems. In this section we provide eigenvalue bounds for the three matrices $K_{1,\text{reg}}$, $K_{2,\text{reg}}$, and $K_{3,\text{reg}}$. We first state known results for the 1×1 and 2×2 block systems, and then move on to present new bounds for the 3×3 block matrix. By continuity of eigenvalues, we choose to start with the regularized formulation and only then move to the unregularized case, because bounds for the latter can be obtained as special cases of the regularized formulation with the regularization parameters set to zero.

4.1. Bounds for the regularized 1×1 and 2×2 block systems. In (2.8c), $K_{1,\text{reg}}$ is positive definite provided that either $\delta > 0$ or J has full row rank. The positive definiteness makes the approach of reducing the original system with the matrix

(2.8a) to the Schur complement system associated with the matrix (2.8c) potentially attractive. On the other hand, reducing the system this way requires the inversion of $H + X^{-1}Z$ and may cause computational difficulties, such as potential loss of sparsity.

A straightforward eigenvalue analysis yields the following result.

THEOREM 4.1 (bounds for $K_{1,\text{reg}}$). *Let $K_{1,\text{reg}}$ be defined as in* (2.8c) *and suppose Assumptions* 2.1 *(convexity) and* 2.2 *(positivity) hold. The eigenvalues of $K_{1,\text{reg}}$ are contained in the interval*

$$\left[\frac{\sigma_m^2}{\lambda_{\max}(H + X^{-1}Z + \rho I)} + \delta,\ \frac{\sigma_1^2}{\lambda_{\min}(H + X^{-1}Z + \rho I)} + \delta\right].$$

As a consequence, we have the following bound on the spectral condition number:

$$\kappa(K_{1,\text{reg}}) \leq \frac{\sigma_1^2 + \delta\lambda_{\min}(H + X^{-1}Z + \rho I)}{\sigma_m^2 + \delta\lambda_{\max}(H + X^{-1}Z + \rho I)}\ \kappa(H + X^{-1}Z + \rho I).$$

No clear relation readily emerges from the bound on the condition number given in Theorem 4.1. However, it is possible to see that when both ρ and δ are positive, the condition number is strictly smaller than that of the unregularized matrix where $\rho = \delta = 0$.

Gondzio (2012) provides similar bounds on the eigenvalues except they are further simplified (and loosened) by replacing the extremal eigenvalues of $H + X^{-1}Z + \rho I$ with a sum of the extremal eigenvalues of H and the maximum or the minimum of $\{z_i/x_i \mid i = 1, \ldots, n\}$. The eigenvalues of $K_{1,\text{reg}}$ are thus contained in the interval

$$\left[\frac{\sigma_m^2}{\lambda_1 + \max(z_i/x_i) + \rho} + \delta,\ \frac{\sigma_1^2}{\lambda_n + \min(z_i/x_i) + \rho} + \delta\right].$$

These looser bounds have the advantage of relying only on the eigenvalues of H, which do not change throughout the interior-point iteration. In contrast, the bounds of Theorem 4.1 are tighter but require the computation of eigenvalues of a different matrix at every interior-point iteration. Asymptotically, in the scenario of Example 2.1, the bound on the condition number reduces to

$$\kappa(K_{1,\text{reg}}) = O(\sigma_1^2/(\rho\delta)).$$

In practice, ρ and δ are allowed to take values as small as about $\sqrt{\varepsilon_{\text{mach}}}$. In this case, there appears to be a definite disadvantage to using the Schur complement equations because the condition number likely exceeds the inverse of machine precision early. Indeed, the implementation of Friedlander and Orban (2012) initializes $\rho = \delta = 1$ and divides both by 10 at each iteration. In IEEE double precision, after just 8 iterations the smallest allowed value of 10^{-8} is reached but convergence is typically not yet achieved.

Results for the regularized K_2 are given in (Friedlander and Orban, 2012, Theorem 5.1), which is itself a specialization of results of Rusten and Winther (1992) and Silvester and Wathen (1994). In the analysis it is not assumed that J has full rank, only that $\rho > 0$ and $\delta > 0$.

THEOREM 4.2 (bounds for $K_{2,\mathrm{reg}}$, Friedlander and Orban (2012)). *Suppose Assumptions* 2.1 *(convexity) and* 2.2 *(positivity) hold, and let* $H_\rho := H + X^{-1}Z + \rho I$. *Then,*

$$\gamma^+ \geq \lambda_{\min}(H_\rho),$$
$$\gamma^+ \leq \tfrac{1}{2}\left((\lambda_{\max}(H_\rho) - \delta) + \sqrt{(\lambda_{\max}(H_\rho) - \delta)^2 + 4(\sigma_1^2 + \delta\lambda_{\max}(H_\rho))}\right)$$

for any positive eigenvalue γ^+ *of* $K_{2,\mathrm{reg}}$, *and*

$$\gamma^- \geq \tfrac{1}{2}\left((\lambda_{\min}(H_\rho) - \delta) - \sqrt{(\lambda_{\min}(H_\rho) - \delta)^2 + 4(\sigma_1^2 + \delta\lambda_{\min}(H_\rho))}\right),$$
$$\gamma^- \leq \tfrac{1}{2}\left((\lambda_{\max}(H_\rho) - \delta) - \sqrt{(\lambda_{\max}(H_\rho) - \delta)^2 + 4(\sigma_m^2 + \delta\lambda_{\max}(H_\rho))}\right)$$

for any negative eigenvalue γ^- *of* $K_{2,\mathrm{reg}}$. *In addition,* $-\delta$ *is an eigenvalue of* $K_{2,\mathrm{reg}}$ *if and only if* J *is rank deficient.*

Again, similar bounds that are looser yet may require fewer computations were derived by Gondzio (2012), and can be obtained by replacing the eigenvalues of H_ρ with the extremal eigenvalues of H and the extremal values of $X^{-1}Z$.

Note that whether or not J is rank deficient, Theorem 4.2 implies that all negative eigenvalues of $K_{2,\mathrm{reg}}$ are bounded above by $-\delta$. Similarly, all positive eigenvalues are bounded below by ρ. The most noticeable effect of regularization in this case is to buffer the eigenvalues away from zero.

In the scenario of Example 2.1, Theorem 4.2 yields the following asymptotic bounds:

$$-\sigma_1 \lesssim \gamma^- \leq -\delta < 0 < \rho \leq \gamma^+ \lesssim \lambda_{\max}(H_\rho) \leq \lambda_1 + \rho + \max_i \frac{z_i}{x_i},$$

so that we obtain the following asymptotic condition number estimate:

$$\kappa(K_{2,\mathrm{reg}}) = O(\lambda_{\max}(H_\rho)/\min(\rho,\delta)) = O(1/(\mu\min(\rho,\delta))).$$

The limits of machine precision, given the common bounds on δ and ρ, are thus not achieved until μ reaches $\sqrt{\varepsilon_{\mathrm{mach}}}$, which typically occurs in the last few iterations.

4.2. Bounds for the regularized 3×3 block system. We now consider the unreduced 3×3 block system. We focus on the symmetrized matrix $\hat{K}_{3,\mathrm{reg}}$ since it allows us to seamlessly work with real arithmetic. We note that the eigenvalues can give valuable information for both symmetric and nonsymmetric matrices. The convergence of minimum residual methods for diagonalizable matrices, for example, depends strongly on the distribution of eigenvalues and the conditioning of the eigenvectors (Saad, 2003, Proposition 6.32). We also note that in our numerical experiments the computed eigenvalues of $\hat{K}_{3,\mathrm{reg}}$ and $K_{3,\mathrm{reg}}$ were nearly identical.

A challenge here is that the 3×3 block form gives rise to rather complicated cubic inequalities in some cases. As we show, simplifying assumptions using the limiting behavior of elements in X and Z lead to effective bounds, although the case of an upper bound on the negative eigenvalues proves to be significantly harder to deal with. The analysis for the 3×3 block system forms the core of our new results. Our technique largely relies on energy estimates in the spirit of Rusten and Winther (1992). We will

be using extensively the fact that for any vectors $u \in \mathbb{R}^n$ and $v \in \text{Null}(J)^\perp \subset \mathbb{R}^n$,

$$\lambda_n \|u\|^2 \leq u^T H u \leq \lambda_1 \|u\|^2, \tag{4.1}$$

$$\sigma_m \|v\| \ \leq \|Jv\| \ \leq \sigma_1 \|v\|. \tag{4.2}$$

Note that the right inequality in (4.2) is satisfied for all $v \in \mathbb{R}^n$.

The eigenvalue problem for $\hat{K}_{3,\text{reg}}$ is formulated as

$$\begin{bmatrix} H + \rho I & J^T & -Z^{\frac{1}{2}} \\ J & -\delta I & 0 \\ -Z^{\frac{1}{2}} & 0 & -X \end{bmatrix} \begin{bmatrix} u \\ v \\ w \end{bmatrix} = \theta \begin{bmatrix} u \\ v \\ w \end{bmatrix}. \tag{4.3}$$

Our first result provides upper and lower bounds on the positive eigenvalues of $\hat{K}_{3,\text{reg}}$. Most notably, we show that upon regularization the lower bound is approximately additively shifted by ρ.

THEOREM 4.3 (bounds on the positive eigenvalues of $\hat{K}_{3,\text{reg}}$). *Suppose Assumption* 2.1 *(convexity) is satisfied. As long as Assumption* 2.2 *(positivity) is satisfied, the positive eigenvalues of $\hat{K}_{3,reg}$ are bounded in $I_+ = [\xi, \eta]$, where*

$$\xi = \min_j \tfrac{1}{2} \left(\lambda_n + \rho - x_j + \sqrt{(\lambda_n + \rho + x_j)^2 + 4 z_j} \right)$$

and η is the largest root of the cubic

$$q_3(\theta) := \theta^3 + (\delta - (\lambda_1 + \rho))\theta^2 - (\delta(\lambda_1 + \rho) + \sigma_1^2 + z_{\max})\theta - z_{\max}\delta.$$

When Assumption 2.3 *(complementarity) holds, ξ reduces to the uniform lower bound $\lambda_n + \rho$.*

Proof. We examine the upper and lower bounds separately.

Upper bound on the positive eigenvalues. We solve for w in the third block row of (4.3) to get $w = -(\theta I + X)^{-1} Z^{\frac{1}{2}} u$, which we substitute into the first block row to obtain

$$(H + \rho I)u + J^T v + Z^{\frac{1}{2}} (\theta I + X)^{-1} Z^{\frac{1}{2}} u = \theta u.$$

The matrices $Z^{\frac{1}{2}}$ and $(\theta I + X)^{-1}$ are diagonal, and therefore commute. Taking the inner product with u gives the following equation for θ:

$$u^T (H + \rho I)u + u^T J^T v + u^T (\theta I + X)^{-1} Z u = \theta \|u\|^2. \tag{4.4}$$

Solving for v in the second block row of (4.3) gives $v = \frac{1}{\theta + \delta} J u$, which we substitute into (4.4) to get

$$u^T (H + \rho I)u + \frac{1}{\theta + \delta} \|Ju\|^2 + u^T (\theta I + X)^{-1} Z u = \theta \|u\|^2. \tag{4.5}$$

We use (4.1) and (4.2) to bound the first and second terms in (4.5):

$$(\lambda_1 + \rho)\|u\|^2 + \frac{\sigma_1^2}{\theta + \delta} \|u\|^2 + u^T (\theta I + X)^{-1} Z u \geq \theta \|u\|^2.$$

Since $(\theta I + X)^{-1}Z$ is diagonal, we have

$$u^T(\theta I + X)^{-1}Zu \le \max_i \frac{z_i}{\theta + x_i}\|u\|^2.$$

We must have $u \neq 0$, since if $u = 0$ then the second block row of (4.3) implies that $(\theta+\delta)v = 0$. Since the matrix (2.8a) is nonsingular, $\theta > 0$ and thus $\theta+\delta > 0$, implying that $v = 0$. Then, by the first block row of (4.3), $Z^{\frac{1}{2}}w = 0$ would imply $w = 0$, which must not occur since an eigenvector must be nonzero. We can thus divide by $\|u\|^2$, and using the last two displayed inequalities we get the relation

$$\lambda_1 + \rho + \frac{\sigma_1^2}{\theta + \delta} + \max_i \frac{z_i}{\theta + x_i} \ge \theta.$$

We can bound the maximum term above by $\frac{z_{\max}}{\theta}$:

$$\lambda_1 + \rho + \frac{1}{\theta + \delta}\sigma_1^2 + \frac{z_{\max}}{\theta} \ge \theta.$$

Multiplying by $(\theta + \delta)\theta$ and rearranging gives

$$\theta^3 + (\delta - (\lambda_1 + \rho))\,\theta^2 - \left(\delta(\lambda_1 + \rho) + \sigma_1^2 + z_{\max}\right)\theta - z_{\max}\delta \le 0.$$

As a consequence, θ must be bounded above by the largest real root of the above cubic polynomial. Note that there is exactly one positive real root, since the values of the cubic and its derivative are negative at zero.

Lower bound on the positive eigenvalues. Taking the inner product of v with the second block row of (4.3) and rearranging, we have

$$v^T Ju = (\theta + \delta)\|v\|^2,$$

which we substitute into (4.4) to give

$$u^T(H + \rho I)u + (\theta + \delta)\|v\|^2 + u^T(\theta I + X)^{-1}Zu = \theta\|u\|^2.$$

Then using (4.1), we have

$$(\lambda_n + \rho)\|u\|^2 + (\theta + \delta)\|v\|^2 + u^T(\theta I + X)^{-1}Zu \le \theta\|u\|^2.$$

The last term on the left-hand side may be bounded below with a minimum:

$$(\lambda_n + \rho)\|u\|^2 + (\theta + \delta)\|v\|^2 + \min_i \frac{z_i}{\theta + x_i}\|u\|^2 \le \theta\|u\|^2. \tag{4.6}$$

We denote the index where the minimum occurs by j, then multiply by $\theta + x_j$ and rearrange into

$$(\theta^2 + (x_j - \lambda_n - \rho)\theta - (x_j(\lambda_n + \rho) + z_j))\|u\|^2 \ge (\theta + \delta)(\theta + x_j)\|v\|^2 \ge 0.$$

Since again $u \neq 0$, we then bound by the positive root of the quadratic. Taking the minimum over all indices j gives the desired bound.

If Assumption 2.3 (complementarity) holds, we can derive a uniform lower bound using the fact that $\min_i z_i/(\theta + x_i) \ge 0$ in (4.6) and obtain

$$(\theta - \lambda_n - \rho)\|u\|^2 \ge (\theta + \delta)\|v\|^2 \ge 0.$$

The lower bound becomes $\theta \ge \lambda_n + \rho$. ∎

Note that in the special case where all z_i converge to 0, and thus $z_{\max} = 0$ at the limit, the upper bound at the limit can be written as an explicit quadratic root independent of z_i:

$$\theta \leq \frac{1}{2}\left(\lambda_1 + \rho - \delta + \sqrt{(\lambda_1 + \rho + \delta)^2 + 4\sigma_1^2}\right).$$

Note in addition that the expression ξ of Theorem 4.3 is always at least equal to $\lambda_n + \rho$. Thus the uniform bound also applies in the course of the iterations.

We begin our investigation of negative eigenvalues with an upper bound, which turns out to depend on the scaling of the problem.

THEOREM 4.4 (an upper bound on the negative eigenvalues of $\hat{K}_{3,\text{reg}}$ during iterations). *Let $\rho \geq 0$, $\delta > 0$, and Assumption 2.1 (convexity) be satisfied. Suppose also that Assumption 2.2 (positivity) holds at (x, y, z), where $x_i > \delta$ for $i = 1, \ldots, n$. Then the negative eigenvalues of $\hat{K}_{3,\text{reg}}$ are bounded above by $-\delta$, and $\theta = -\delta$ is an eigenvalue if and only if J is rank deficient.*

Proof. We first show that $-\delta$ is an upper bound. Assume by contradiction that there is a negative eigenvalue that satisfies $\theta > -\delta$. Upon extracting $v = \frac{1}{\theta+\delta}Ju$ from the second block row of (4.3) and using the identity $w^T Z^{\frac{1}{2}} u = -w^T(\theta I + X)w$ from the third block row, taking the inner product of the first block row with u gives

$$u^T(H + \rho I)u + \frac{1}{\theta + \delta}\|Ju\|^2 + w^T(\theta I + X)w = \theta\|u\|^2.$$

Since $\theta + \delta > 0$ by assumption and all $x_i > \delta > -\theta$, the left-hand side of the last identity is positive. If $u = 0$, then both v and w are also zero, giving a trivial eigenvector and therefore a contradiction. If $u \neq 0$, θ must be positive, which contradicts our initial assumption on the sign of θ. Thus the negative eigenvalues are bounded above by $-\delta$.

We now move on to show when $-\delta$ is an eigenvalue. If J is rank deficient, then $u = 0$, $0 \neq v \in \text{Null}(J^T)$, and $w = 0$ satisfies (4.3) with $\theta = -\delta$. Suppose now that J has full rank. We will show that $\theta \neq -\delta$. By contradiction, assume that $\theta = -\delta$. From the third block row and the assumption that all $x_i > \delta$, we have

$$w^T Z^{\frac{1}{2}} u = w^T(\delta I - X)w \leq 0.$$

Taking the inner product of the first block row of (4.3) with u and using the above inequality and $Ju = 0$ from the second block row, we obtain

$$-\delta\|u\|^2 = u^T(H + \rho I)u - u^T Z^{\frac{1}{2}} w \geq 0.$$

Since $\delta > 0$, this must mean that $u = 0$. The third block row of (4.3) then gives $w = 0$ and we are left with $J^T v = 0$ in the first block row. Since J has full row rank, we conclude that $v = 0$ and that $\theta = -\delta$ cannot be an eigenvalue. ∎

Interestingly, a similar result holds in the limit, as we show below. We note however that there seems to be a transition zone between the moment when some components of x drop below δ and the limit when strict complementarity applies. This "gray zone" is necessarily attained if $\mathcal{A} \neq \emptyset$, and it is more difficult to characterize the relationship between θ and $-\delta$ in that zone. Theorem 4.5 below assumes that the optimization problem has been scaled appropriately prior to solving. The "gray zone"

is strongly tied to the quality of our scaling assumptions; better scaling may shrink that zone.

THEOREM 4.5 (an upper bound on the negative eigenvalues of $\hat{K}_{3,\text{reg}}$ in the limit). *Let $\rho > 0$, $\delta > 0$, and Assumption 2.1 (convexity) be satisfied. Suppose also Assumption 2.3 (complementarity) holds at (x, y, z), where strict complementarity is satisfied, $x_i > \delta$ for all $i \in \mathcal{I}$, and $\max_i \sqrt{z_i}$ is sufficiently small. Then the negative eigenvalues of $\hat{K}_{3,\text{reg}}$ are bounded above by $-\delta$, and $\theta = -\delta$ is an eigenvalue if and only if J is rank deficient.*

Proof. We first show that $-\delta$ is an upper bound on the negative eigenvalues. Assume by contradiction that there exists a negative eigenvalue that satisfies $\theta > -\delta$. Since $\hat{K}_{3,\text{reg}}$ is nonsingular, there must exist an $\epsilon > 0$ such that $\theta \leq -\epsilon$ for all negative eigenvalues. If $\delta \leq \epsilon$ this implies $-\epsilon \leq -\delta$ and the eigenvalues are bounded above by $-\delta$, which would be in line with the statement of the theorem. So let us assume that $\epsilon < \delta$. In the limit, we have $x_{\mathcal{A}} = 0$, $x_{\mathcal{I}} > 0$, and $z_{\mathcal{I}} = 0$. Because strict complementarity is satisfied, we must also have $z_{\mathcal{A}} > 0$. Partitioning the third block row of $\hat{K}_{3,\text{reg}}$ in (4.3) into active and inactive components gives

$$-Z_{\mathcal{A}\mathcal{A}}^{\frac{1}{2}} u_{\mathcal{A}} = \theta w_{\mathcal{A}}, \tag{4.7}$$

$$-X_{\mathcal{I}\mathcal{I}} w_{\mathcal{I}} = \theta w_{\mathcal{I}}. \tag{4.8}$$

We may rewrite (4.8) as $(X_{\mathcal{I}\mathcal{I}} + \theta I) w_{\mathcal{I}} = 0$, which implies $w_{\mathcal{I}} = 0$ because $x_i + \theta > x_i - \delta > 0$ for all $i \in \mathcal{I}$ by assumption. Taking the inner product of both sides of (4.7) with $w_{\mathcal{A}}$ gives

$$-w_{\mathcal{A}}^T Z_{\mathcal{A}\mathcal{A}}^{\frac{1}{2}} u_{\mathcal{A}} = \theta \|w_{\mathcal{A}}\|^2. \tag{4.9}$$

Taking now the inner product of the first block row of (4.3) with u, the inner product of the second block row with v, and combining them, we write

$$\theta \|u_{\mathcal{I}}\|^2 = u^T (H + \rho I) u + (\theta + \delta) \|v\|^2 + \theta(\|w_{\mathcal{A}}\|^2 - \|u_{\mathcal{A}}\|^2), \tag{4.10}$$

where we used the decomposition $\|u\|^2 = \|u_{\mathcal{A}}\|^2 + \|u_{\mathcal{I}}\|^2$ and (4.9). Note that from (4.7), $u_{\mathcal{A}} = 0$ if and only if $w_{\mathcal{A}} = 0$. If both vanish, necessarily $u_{\mathcal{I}} \neq 0$, and then the right-hand side of (4.10) is strictly positive. This would imply that $\theta > 0$, a contradiction. By our assumption that all $\sqrt{z_i}$ are sufficiently small, we suppose from this point on that $\max_i \sqrt{z_i} < \epsilon$. Suppose now that $u_{\mathcal{A}} \neq 0$ and $w_{\mathcal{A}} \neq 0$. Rearranging (4.7) we find that $w_{\mathcal{A}} = -\frac{1}{\theta} Z_{\mathcal{A}\mathcal{A}}^{\frac{1}{2}} u_{\mathcal{A}}$, and using the upper bounds $\max_i \sqrt{z_i} < \epsilon$ and $\theta \leq -\epsilon$ we have $\|w_{\mathcal{A}}\|^2 \leq \frac{\epsilon^2}{\theta^2} \|u_{\mathcal{A}}\|^2 \leq \|u_{\mathcal{A}}\|^2$. Substituting into (4.10) gives

$$\theta \|u_{\mathcal{I}}\|^2 \geq u^T (H + \rho I) u + (\theta + \delta) \|v\|^2.$$

The right-hand side above is strictly positive, and if $u_{\mathcal{I}} \neq 0$ we have $\theta > 0$, a contradiction. If $u_{\mathcal{I}} = 0$, then $u_{\mathcal{A}} = 0$, again a contradiction. Therefore, we cannot have $\theta > -\delta$, and we conclude that $\theta \leq -\delta$.

We now move on to find when $-\delta$ is an eigenvalue. If J is rank deficient, then as in Theorem 4.4, $(0, v, 0)$ with $0 \neq v \in \text{Null}(J^T)$ is an eigenvector for $\theta = -\delta$. Now suppose that J has full rank, and assume by contradiction that $\theta = -\delta$. Partitioning as above gives

$$Z_{\mathcal{A}\mathcal{A}}^{\frac{1}{2}} u_{\mathcal{A}} = \delta w_{\mathcal{A}}, \tag{4.11}$$

$$X_{\mathcal{I}\mathcal{I}} w_{\mathcal{I}} = \delta w_{\mathcal{I}}. \tag{4.12}$$

Upon rearranging (4.12), the matrix $(X_{\mathcal{I}\mathcal{I}} - \delta I)$ has full rank by assumption, so we have $w_{\mathcal{I}} = 0$. If $u = 0$ we also have $w_{\mathcal{A}} = 0$ and there only remains $J^T v = 0$ in the first block row of (4.3), giving a contradiction because we obtain a trivial solution for the eigenvector. Thus $u \neq 0$. Taking the inner product of (4.11) with $w_{\mathcal{A}}$ reveals that

$$w^T Z^{\frac{1}{2}} u = w_{\mathcal{A}}^T Z_{\mathcal{A}\mathcal{A}}^{\frac{1}{2}} u_{\mathcal{A}} = \delta \|w_{\mathcal{A}}\|^2. \tag{4.13}$$

Using the Cauchy–Schwarz inequality and the bound on $\max_i \sqrt{z_i}$, we have

$$w_{\mathcal{A}}^T Z_{\mathcal{A}\mathcal{A}}^{\frac{1}{2}} u_{\mathcal{A}} \leq \epsilon \|w_{\mathcal{A}}\| \|u_{\mathcal{A}}\|,$$

and using (4.13) and rearranging gives

$$\|w_{\mathcal{A}}\| \leq \frac{\epsilon}{\delta} \|u_{\mathcal{A}}\| \leq \|u_{\mathcal{A}}\| \leq \|u\|, \tag{4.14}$$

since $\epsilon < \delta$. Taking the inner product of the first block row with u, using (4.13) and $Ju = 0$ from the second block row, and rearranging, we have

$$u^T (H + \rho I) u = \delta(\|w_{\mathcal{A}}\|^2 - \|u\|^2),$$

which reduces to $u^T(H + \rho I)u \leq 0$ by (4.14). Therefore $u = 0$, a contradiction, and hence $\theta \neq -\delta$ when J is full rank. ∎

Next, we derive a lower bound on the negative eigenvalues.

THEOREM 4.6 (a lower bound on the negative eigenvalues of $\hat{K}_{3,\text{reg}}$). *Suppose Assumption* 2.1 *(convexity) is satisfied. Assume $\theta I + X$ is nonsingular for all θ throughout the computation. Then the negative eigenvalues θ of the matrix $\hat{K}_{3,\text{reg}}$ satisfy $\theta \geq \zeta$, where*

$$\zeta := \min \left\{ \tfrac{1}{2} \left(\lambda_n + \rho - \delta - \sqrt{(\lambda_n + \rho + \delta)^2 + 4\sigma_1^2} \right), \min_{\theta + x_j < 0} \theta_j^* \right\}, \tag{4.15}$$

and θ_j^* is the smallest root of the cubic

$$\begin{aligned} \bar{q}_3(\theta) := \theta^3 &+ \left(x_j + \delta - \lambda_n - \rho\right) \theta^2 + \left(\delta x_j - \delta(\lambda_n + \rho) - x_j(\lambda_n + \rho) - \sigma_1^2 - z_j\right) \theta \\ &- \left(\delta x_j (\lambda_n + \rho) + \sigma_1^2 x_j + z_j \delta\right). \end{aligned}$$

Proof. We assume that $\theta + \delta < 0$. The case where $\theta \geq -\delta$ poses no difficulty because it is easy to verify that $\zeta \leq -\delta$.

We start from (4.5) with the bounds in (4.1) and (4.2) to get

$$(\lambda_n + \rho)\|u\|^2 + \frac{1}{\theta + \delta} \sigma_1^2 \|u\|^2 + u^T (\theta I + X)^{-1} Z u \leq \theta \|u\|^2.$$

We note that $u \neq 0$, since if $u = 0$ the second block row of (4.3) implies $(\theta + \delta)v = 0$, implying that $v = 0$. The first line yields $Z^{\frac{1}{2}} w = 0$ and thus $w = 0$, a contradiction. Bounding the last term of the left side of the previous inequality by the minimum,

$$\left(\lambda_n + \rho + \frac{1}{\theta + \delta} \sigma_1^2 + \min_i \frac{z_i}{\theta + x_i} \right) \leq \theta.$$

We consider two cases. In case one, $\theta + x_i > 0$ for all indices i, and in this case we can bound the minimum term from below by zero. In case two, some $\theta + x_i < 0$, and there exists an index j such that $\min_i \frac{z_i}{\theta+x_i} = \frac{z_j}{\theta+x_j}$.

Case one. We bound the minimum term below by zero in the inequality above,

$$\lambda_n + \rho + \frac{1}{\theta + \delta}\sigma_1^2 \leq \theta, \tag{4.16}$$

which we can multiply by $(\theta + \delta)$ and rearrange to give

$$\theta^2 + (\delta - \lambda_n - \rho)\theta - (\delta(\lambda_n + \rho) + \sigma_1^2) \leq 0.$$

Therefore,

$$\theta \geq \tfrac{1}{2}\left(\lambda_n + \rho - \delta - \sqrt{(\delta - \lambda_n - \rho)^2 + 4(\delta(\lambda_n + \rho) + \sigma_1^2)}\right).$$

Case two. We use the index j for the minimum to get

$$\lambda_n + \rho + \frac{1}{\theta + \delta}\sigma_1^2 + \frac{z_j}{\theta + x_j} \leq \theta.$$

Multiplying by $(\theta + \delta)(\theta + x_j)$ and rearranging, we get

$$\theta^3 + (x_j + \delta - \lambda_n - \rho)\theta^2 + (\delta x_j - \delta(\lambda_n + \rho) - x_j(\lambda_n + \rho) - \sigma_1^2 - z_j)\theta \\ - (\delta x_j(\lambda_n + \rho) + \sigma_1^2 x_j + z_j\delta) \geq 0.$$

We then define θ_j^* to be the smallest root of the cubic above.

The bound stated in the theorem is given by the minimum of these possible bounds from cases one and two. ∎

We remark that in practice we cannot know which indices j satisfy $\theta + x_j < 0$, but we can simply compute θ_j^* for all indices j and use this in the comparison. Second, we note that the possibility that some $\theta < 0$ in the spectrum of $\hat{K}_3$ be an eigenvalue of $-X$ could arise in the course of the iterations or in the limit if there are inactive bounds.

Consider now the scenario of Example 2.1. Assuming that the problem has been scaled appropriately, the bounds of the previous results simplify to

$$\zeta \leq \theta \leq -\delta < 0 \quad \text{or} \quad 0 < \lambda_n + \rho \leq \theta \leq \eta,$$

where ζ and η are both finite. Thus we obtain the asymptotic condition number estimate

$$\kappa(K_{3,\text{reg}}) \leq \max(\eta, -\zeta)/\min(\rho + \lambda_n, \delta) = O\left(1/\min(\rho + \lambda_n, \delta)\right).$$

Here we have the validation of our claim that the block 3×3 system sees the best conditioning. Under the usual choices of ρ and δ and unless the conditioning of the problem is such that η or ζ is very large, this condition number will remain within computing limits through convergence of the iteration. Our numerical experiments, presented in section 6, verify that the 3×3 matrices remain numerically nonsingular—and reasonably conditioned—throughout.

4.3. Comparisons with existing bounds. It is instructive to ask whether the bounds we provide in this paper are sharper than existing bounds. In this section we compare our results with known bounds from the literature and show that, indeed, for the cases we have been able to analyze, our bounds are generally sharper. This is not surprising because our bounds exploit the specific block structure of our saddle-point matrices.

Silvester and Wathen (1994, Lemma 2.2) generalize the results of Rusten and Winther (1992) for saddle-point matrices to the case of a nonzero $(2,2)$ block. They focus on the discretized Stokes equations, but their paper contains several results that are applicable to general regularized saddle-point systems with a positive definite leading block and a positive semidefinite $(2,2)$ block. In the context of the present paper, we can apply the Silvester and Wathen results by splitting

$$K_{3,\mathrm{reg}} = \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix}, \quad \text{where} \quad A := H + \rho I, \quad B := \begin{bmatrix} J \\ -Z^{1/2} \end{bmatrix}, \quad C := \begin{bmatrix} \delta I & 0 \\ 0 & X \end{bmatrix}.$$

In their analysis, Silvester and Wathen use the extremal eigenvalues of A, the extremal singular values of B, and a bound on the norm of C, which here is $\|C\| \leq \max(\delta, x_{\max})$, where $x_{\max} := \max_{1 \leq i \leq n} x_i$. Note that the singular values of B are the square roots of the eigenvalues of $B^T B = J^T J + Z$.

In this section we denote generic positive and negative eigenvalues of $K_{3,\mathrm{reg}}$ by θ^+ and θ^-, respectively. Silvester and Wathen (1994, Lemma 2.2) provide the upper bound

$$\begin{aligned} \theta^+ &\leq \frac{1}{2}\left(\lambda_{\max}(H+\rho I) + \sqrt{\lambda^2_{\max}(H+\rho I) + 4\lambda_{\max}(J^TJ+Z)}\right) \\ &\leq \frac{1}{2}\left(\lambda_1 + \rho + \sqrt{(\lambda_1+\rho)^2 + 4(\sigma_1^2 + z_{\max})}\right), \end{aligned} \tag{4.17}$$

where we used our notation for the eigenvalues of H and singular values of J, and the facts that $\lambda_{\max}(H+\rho I) = \lambda_1 + \rho$ and $\sigma^2_{\max}(B) \equiv \lambda_{\max}(J^TJ + Z) \leq \sigma_1^2 + z_{\max}$, where $z_{\max} := \max_{1 \leq i \leq n} z_i$. The transition from the first row of (4.17) to the second row is a natural relaxation whose purpose is to make the bounds practical and directly connected to the blocks of $K_{3,\mathrm{reg}}$. Note that contrary to the bound given by Theorem 4.3, the parameter δ plays no role in (4.17).

It is straightforward to show that when $\delta = 0$, the largest positive root of the cubic polynomial q_3 of Theorem 4.3 reduces precisely to (4.17). Denote this root by θ_0. When $\delta > 0$ we claim that the largest root of q_3, denoted θ_δ, decreases, which means that the upper bound on θ^+ becomes tighter. To show this, we draw from the perturbation analysis of Wilkinson (1959, section 3). Let us first write $q_3(\theta) = p_3(\theta) + \delta p_2(\theta)$, where

$$p_3(\theta) := \theta^3 - (\lambda_1+\rho)\theta^2 - (\sigma_1^2 + z_{\max})\theta \qquad \text{and} \qquad p_2 := \theta^2 - (\lambda_1+\rho)\theta - z_{\max}.$$

If the polynomial p_3 is perturbed by δp_2, the root is perturbed according to

$$\theta_\delta = \theta_0 + \frac{\mathrm{d}\theta_0}{\mathrm{d}\delta}\delta + \cdots, \qquad \text{where} \quad \frac{\mathrm{d}\theta_0}{\mathrm{d}\delta} = -\frac{p_2(\theta_0)}{p_3'(\theta_0)},$$

where the derivative is interpreted as the derivative of θ_δ with respect to δ subsequently

evaluated at $\delta = 0$. In order to evaluate $p_2(\theta_0)$, first note that

$$\begin{aligned}\theta_0^2 &= \frac{1}{2}\left[(\lambda_1+\rho)\left(\lambda_1+\rho+\sqrt{(\lambda_1+\rho)^2+4(\sigma_1^2+z_{\max})}\right)+2(\sigma_1^2+z_{\max})\right]\\ &= (\lambda_1+\rho)\theta_0+(\sigma_1^2+z_{\max}).\end{aligned}$$

We conclude that $p_2(\theta_0) = \sigma_1^2 > 0$. On the other hand, it is straightforward to confirm that $p_3'(\theta_0) > 0$ and we conclude that $\mathrm{d}\theta_0/\mathrm{d}\delta < 0$. Therefore, at least for small values of δ, $\theta_\delta < \theta_0$, and hence our bound is sharper than (4.17).

Next, Silvester and Wathen (1994, Lemma 2.2) provide the lower bound

$$\theta^+ \geq \lambda_{\min}(H+\rho I) = \lambda_n + \rho.$$

The latter coincides with the lower bound of Theorem 4.3 in the special case where $z = 0$, which is a rather unlikely scenario. In all other cases, our bound is sharper. This is most easily seen from (4.6), which may be written

$$\theta \geq \lambda_n + \rho + \min_i \frac{z_i}{\theta + x_i} > \lambda_n + \rho$$

for any positive eigenvalue θ, where the last inequality is strict in the course of the iterations and typically holds as an equality in the limit. Note also that from the expression in Theorem 4.3, we have $\xi \geq \lambda_n + \rho$ for any nonnegative values of x_j and z_j, $j = 1, \ldots, n$. In addition, ξ is a decreasing function of x_j and an increasing function of z_j with $\lim_{x_j \to \infty} \xi = \lambda_n + \rho$, where the limit is taken over all indices j.

Regarding negative eigenvalues, Silvester and Wathen (1994, Lemma 2.2) provide the upper bound

$$\begin{aligned}\theta^- &\leq \frac{1}{2}\left(\lambda_{\max}(H+\rho I) - \sqrt{\lambda_{\max}^2(H+\rho I)+4\lambda_{\min}(J^TJ+Z)}\right)\\ &\leq \frac{1}{2}\left(\lambda_1+\rho-\sqrt{(\lambda_1+\rho)^2+4(\sigma_m^2+z_{\min})}\right),\end{aligned}$$

where $z_{\min} = \min_i z_i$. This bound is again independent of δ. When J is rank deficient, $\sigma_m = 0$ and the above bound reduces to $-z_{\min}$, which typically converges to zero. By contrast, under the assumptions of Theorem 4.4, our bound is $-\delta$ and hence is sharper.

Finally, Silvester and Wathen (1994, Lemma 2.2) provide the lower bound

$$\begin{aligned}\theta^- &\geq \frac{1}{2}\left(\lambda_{\min}(H+\rho I) - \sqrt{\lambda_{\min}^2(H+\rho I)+4\lambda_{\max}(J^TJ+Z)}\right)\\ &\geq \frac{1}{2}\left(\lambda_n+\rho-\sqrt{(\lambda_n+\rho)^2+4(\sigma_1^2+z_{\max})}\right). \qquad (4.18)\end{aligned}$$

Again, there is no dependence in δ. This bound is more complicated to analyze and we now show that, except in certain cases, it is tighter than the bound of Theorem 4.6. The value (4.18) is the smallest root of the quadratic

$$p_2(\theta) := \theta^2 - (\lambda_n+\rho)\theta - (\sigma_1^2+z_{\max}).$$

Our lower bound is given by the value ζ of Theorem 4.6. Suppose temporarily that the first term realizes the minimum in (4.15). In this case, ζ is the smallest root of the quadratic

$$\begin{aligned}\bar{p}_2(\theta) &:= \theta^2 - (\lambda_n+\rho-\delta)\theta - (\sigma_1^2+\delta(\lambda_n+\rho))\\ &= p_2(\theta) + \delta(\theta-(\lambda_n+\delta)) + z_{\max}.\end{aligned}$$

Wilkinson's analysis generalizes easily to the case of perturbations of the form above. More precisely, let $\bar{p}(\theta) = p(\theta) + \delta r(\theta) + zs(\theta)$, where δ and z are perturbation parameters, let $\beta = \beta(0,0)$ be a simple root of p, i.e., a simple root of $\bar{p}$ when $\delta = z = 0$, and let $\beta(\delta, z)$ be the corresponding root of $\bar{p}$ for given values of δ and z. Then,

$$\beta(\delta, z) = \beta + \delta \frac{\mathrm{d}\beta}{\mathrm{d}\delta} + z \frac{\mathrm{d}\beta}{\mathrm{d}z} + \cdots, \quad \text{where} \quad \frac{\mathrm{d}\beta}{\mathrm{d}\delta} = -\frac{r(\beta)}{p'(\beta)}, \quad \frac{\mathrm{d}\beta}{\mathrm{d}z} = -\frac{s(\beta)}{p'(\beta)}.$$

More concisely, $\beta(\delta, z) = \beta - (\delta r(\beta) + zs(\beta))/p'(\beta) + \cdots$.

We may apply the above perturbation estimate to our situation by understanding β as the bound of (4.18) and by selecting $\bar{p} = \bar{p}_2$, $p = p_2$, $r(\theta) := \theta - (\lambda_n + \delta)$, and $s(\theta) := 1$ to obtain

$$\zeta = \beta - \frac{\delta(\beta - (\lambda_n + \rho)) + z_{\max}}{p_2'(\beta)} + \cdots.$$

Because $p_2'(\beta) < 0$ and $\beta - (\lambda_n + \rho) < 0$, we have $\zeta > \beta$ provided that

$$\delta > \frac{z_{\max}}{\lambda_n + \rho - \beta} = \frac{2\, z_{\max}}{\lambda_n + \rho + \sqrt{(\lambda_n + \rho)^2 + 4(\sigma_1^2 + z_{\max})}}.$$

In particular, our bound is sharper in the unlikely scenario where $z_{\max} = 0$ provided $\delta > 0$. In all other cases, our bound is tighter provided $z_{\max}$ is sufficiently small.

We now turn our attention to the case where ζ is the second term in the minimum of the statement of Theorem 4.6, i.e., the smallest root of the cubic

$$\bar{q}_3(\theta) = p_3(\theta) + \delta q_2(\theta) + x_j \bar{q}_2(\theta) + \delta x_j q_1(\theta),$$

where

$$\begin{aligned} p_3(\theta) &:= \theta\left(\theta^2 - (\lambda_n + \rho)\theta - (\sigma_1^2 + z_j)\right), & q_2(\theta) &:= \theta^2 - (\lambda_n + \rho)\theta - z_j, \\ \bar{q}_2(\theta) &:= \theta^2 - (\lambda_n + \rho)\theta + \sigma_1^2, & q_1(\theta) &:= \theta - (\lambda_n + \rho). \end{aligned}$$

The bound β given by Silvester and Wathen is the smallest root of $p_3(\theta)$ when $z_j = z_{\max}$, the latter being the index that minimizes the value of β over all $j = 1, \ldots, n$. Let us thus assume that $z_j = z_{\max}$ and the corresponding x_j. Proceeding as above, we have, to first order,

$$\zeta = \beta - \frac{\delta q_2(\beta) + x_j \bar{q}_2(\beta) + \delta x_j q_1(\beta)}{p_3'(\beta)} + \cdots.$$

It is not difficult to verify that $\beta^2 = (\lambda_n + \rho)\beta + \sigma_1^2 + z_{\max}$ and therefore that $q_2(\beta) = \sigma_1^2 > 0$, $\bar{q}_2(\beta) = 2\sigma_1^2 + z_{\max} > 0$, and $q_1(\beta) < 0$. In addition, $p_3'(\beta) = \beta(2\beta - (\lambda_n + \rho)) > 0$. These inequalities combine with the above first-order expansion to establish that $\zeta > \beta$, i.e., the bound of Theorem 4.6 is tighter than (4.18), whenever

$$\tfrac{1}{2}\delta x_j \left(\lambda_n + \rho + \sqrt{(\lambda_n + \rho)^2 + 4(\sigma_1^2 + z_{\max})}\right) > \delta\sigma_1^2 + x_j(2\sigma_1^2 + z_{\max}).$$

We now consider more general results introduced by Bai (2013), who derives bounds for saddle-point matrices with an indefinite leading block. In this setting, we write

$$K_{3,\mathrm{reg}} = \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix}, \quad \text{where} \quad A := \begin{bmatrix} H + \rho I & J^T \\ J & -\delta I \end{bmatrix}, \; B := \begin{bmatrix} -Z^{1/2} & 0 \end{bmatrix}, \; C := X.$$

Note that because A is quasi definite, A^{-1} is also quasi definite with the same block structure (Vanderbei, 1995). In particular, the leading block of A^{-1} is positive definite. The Schur complement of A in $K_{3,\text{reg}}$ is

$$-X - \begin{bmatrix} Z^{1/2} & 0 \end{bmatrix} \begin{bmatrix} H + \rho I & J^T \\ J & -\delta I \end{bmatrix}^{-1} \begin{bmatrix} Z^{1/2} \\ 0 \end{bmatrix}$$

and is therefore negative definite. Considering this fact, the result that is relevant to the present situation is (Bai, 2013, Theorem 3.2 (i)). In Bai's notation, let $\Psi > 0$ be an upper bound on the eigenvalues of $BA^{-2}B^T$. This upper bound is used to define two constants $\phi \in (0, 1)$ and $\Phi > 1$. Bai's theorem 3.2 (i) then states that the positive eigenvalues of $K_{3,\text{reg}}$ are located in the interval $[\phi\lambda^+_{\min}(A), \Phi\lambda_{\max}(A)]$, where $\lambda^+_{\min}(A)$ is the smallest positive eigenvalue of A. It is now easy to see that Bai's lower and upper bounds are Silvester and Wathen's bounds multiplied by factors $\phi \in (0, 1)$ and $\Phi > 1$, respectively, and are therefore looser. It must be stressed that Bai's bounds apply to a substantially more general case, for which they are effective.

It is difficult to compare our bounds on the negative eigenvalues to the corresponding result in (Bai, 2013, Theorem 3.2 (i)), because the latter relies on bounds on the eigenvalues of the Schur complement. In our setting, assuming such knowledge is impractical.

5. Eigenvalue bounds for the unregularized systems. Eigenvalue analysis for systems (2.2), (2.3), and (2.4) can be straightforwardly done as special cases of the analysis of section 4. We choose to state these results separately, as corollaries, due to the importance of the unregularized approach in the implementation of interior-point solvers.

5.1. Bounds for the unregularized 1×1 and 2×2 block systems. For the Schur complement system, the following result is a straightforward consequence of Theorem 4.1.

COROLLARY 5.1 (bounds for K_1). *Let K_1 be defined as in* (2.4) *and suppose Assumptions* 2.1 *(convexity) and* 2.2 *(positivity) hold and J has full row rank. The eigenvalues of K_1 are contained in the interval*

$$\left[\frac{\sigma_m^2}{\lambda_{\max}(H + X^{-1}Z)}, \ \frac{\sigma_1^2}{\lambda_{\min}(H + X^{-1}Z)} \right].$$

As a consequence, we have an upper bound on the spectral condition number of K_1:

$$\kappa(K_1) \leq \kappa(J)^2 \ \kappa(H + X^{-1}Z).$$

In the typical situation of Example 2.1, it is clear that K_1 approaches singularity. We have the asymptotic estimates

$$\lambda_{\min}(H + X^{-1}Z) = \Omega(\lambda_n + \mu) \quad \text{and} \quad \lambda_{\max}(H + X^{-1}Z) = \Theta(1/\mu). \tag{5.1}$$

As $\mu \to 0$, Corollary 5.1 thus yields the asymptotic estimate

$$\kappa(K_1) = \kappa(J)^2 \, O\Big(1/(\mu(\lambda_n + \mu))\Big). \tag{5.2}$$

Note that at least in the case of linear programming, the bounds of Corollary 5.1 and (5.2) are tight and (5.2) becomes $\kappa(K_1) = \kappa(J)^2 \, \Theta(1/\mu^2)$. On the other hand, when

$\lambda_n > 0$, (5.2) becomes $\kappa(K_1) = \kappa(J)^2\, O(1/\mu)$. This last estimate is in line with those of Wright (1994), who assumes that a second-order sufficiency condition holds.

Similarly to the situation of Theorem 4.1, bounds in the spirit of Gondzio (2012) can be obtained by substituting $\rho = \delta = 0$. The eigenvalues can thus be bounded by expressions involving only the extremal eigenvalues of H and singular values of J. Asymptotic bounds turn out to be identical in terms of order of magnitude to the ones based on Corollary 5.1.

Moving on to consider the 2×2 block system, Corollary 3.6 states that K_2 is nonsingular during the iterations if and only if J has full row rank. For this saddle-point linear system much is known in the literature, and we state now observations that can be concluded from existing results, e.g., (Rusten and Winther, 1992, section 2).

The following result readily follows from Theorem 4.2, but can in fact be directly obtained from (Rusten and Winther, 1992, Lemma 2.1).

COROLLARY 5.2 (bounds for K_2). *If H is positive semidefinite, J has full row rank and Assumptions* 2.1 *(convexity) and* 2.2 *(positivity) are satisfied, then*

$$\gamma^+ \geq \lambda_{\min}(H + X^{-1}Z),$$
$$\gamma^+ \leq \tfrac{1}{2}\left(\lambda_{\max}(H + X^{-1}Z) + \sqrt{\lambda_{\max}(H + X^{-1}Z)^2 + 4\sigma_1^2}\right)$$

for any positive eigenvalue γ^+ of K_2, and

$$\gamma^- \geq \tfrac{1}{2}\left(\lambda_{\min}(H + X^{-1}Z) - \sqrt{\lambda_{\min}(H + X^{-1}Z)^2 + 4\sigma_1^2}\right),$$
$$\gamma^- \leq \tfrac{1}{2}\left(\lambda_{\max}(H + X^{-1}Z) - \sqrt{\lambda_{\max}(H + X^{-1}Z)^2 + 4\sigma_m^2}\right)$$

for any negative eigenvalue γ^- of K_2.

From Corollary 5.2 we see that the lower bound on the negative eigenvalues of K_2 is finite and bounded away from zero unless $J = 0$. It is the other three bounds that are responsible for the ill-conditioning of K_2. Using again the situation of Example 2.1 where the extremal eigenvalues of $H + X^{-1}Z$ are approximated by (5.1), we obtain the asymptotic estimates

$$\lambda_n + \mu \lesssim \gamma^+ \lesssim 1/\mu \qquad \text{and} \qquad \tfrac{1}{2}\left(\lambda_n - \sqrt{\lambda_n^2 + 4\sigma_1^2}\right) \lesssim \gamma^- \lesssim -\mu\,\sigma_m^2,$$

where the upper bound on the negative eigenvalues is obtained via the Taylor expansion $\sqrt{1+x} \approx 1 + \frac{1}{2}x$ for $x \ll 1$. These estimates yield the asymptotic condition number

$$\kappa(K_2) = O\left(\frac{1/\mu}{\min(\mu\sigma_m^2, \lambda_n + \mu)}\right) = O\left(\frac{1}{\mu^2}\right).$$

Several authors, including Fourer and Mehrotra (1993) and Korzak (1999) suggest scalings of K_2 that alleviate this ill-conditioning.

The above asymptotic estimates must be considered cautiously, as they do not always fully capture the actual value of the condition number. Nevertheless, they illustrate the ill-conditioning of the 1×1 and 2×2 formulations in the unregularized case.

5.2. Bounds for the unregularized 3×3 block system. We now perform a similar eigenvalue analysis for the matrix $\hat{K}_3$. Here again, our results can be obtained as special cases of the analysis of section 4. Moulding (2012) provides detailed and direct proofs that are derived independently of the regularized case.

COROLLARY 5.3 (bounds on the positive eigenvalues of $\hat{K}_3$). *Suppose Assumption* 2.1 *(convexity) holds. For as long as Assumption* 2.2 *(positivity) holds and J has full rank, the positive eigenvalues of $\hat{K}_3$ are bounded in*

$$I_+ := \left[\min_j \tfrac{1}{2} \left(\lambda_n - x_j + \sqrt{(\lambda_n + x_j)^2 + 4z_j} \right), \ \tfrac{1}{2} \left(\lambda_1 + \sqrt{\lambda_1^2 + 4(\sigma_1^2 + z_{\max})} \right) \right],$$

where $z_{\max} := \max_i z_i$. When Assumption 2.3 *(complementarity) holds, the lower bound reduces to $\lambda_n \geq 0$.*

The proof follows from Theorem 4.3 with $\delta = \rho = 0$, but observe that we see a significant simplification here for the upper bound. This is because for $\delta = 0$ the cubic equation at the statement of Theorem 4.3 simplifies to the quadratic equation

$$\theta^2 - \lambda_1 \theta - (\sigma_1^2 + z_{\max}) = 0.$$

Note that the lower bound on positive eigenvalues in Corollary 5.3 is strictly larger than λ_n as long as Assumption 2.2 (positivity) holds. In a sufficiently small neighborhood of an isolated minimizer, the minimum term in the lower bound will be attained for some $j \in \mathcal{I}$. This lower bound is strictly positive as long as $(x, z) > 0$ but in the limit, by definition of $\mathcal{I}$, it reduces to λ_n. If $\lambda_n = 0$, this limiting bound is overly pessimistic and is not tight if the LICQ is satisfied since $\hat{K}_3$ is nonsingular by the results in section 3.3.

Next, we consider bounds on the negative eigenvalues. We are only able to find an effective lower bound; the upper bound that we find is zero and is not particularly meaningful. The corollary below follows from Theorem 4.6, which applies also to zero values of the regularization parameters, namely, $\delta = \rho = 0$.

COROLLARY 5.4 (bounds on the negative eigenvalues of $\hat{K}_3$). *Suppose Assumption* 2.1 *(convexity) holds. Suppose the matrix $\theta I + X$ is nonsingular for all $\theta < 0$ in the spectrum of $\hat{K}_3$. The negative eigenvalues of $\hat{K}_3$ are bounded in $I_- = [\zeta, 0)$, where*

$$\zeta := \min \left\{ \tfrac{1}{2} \left(\lambda_n - \sqrt{\lambda_n^2 + 4\sigma_1^2} \right), \min_{\{j | \theta + x_j < 0\}} \theta_j^* \right\}$$

and θ_j^* is the smallest root of the cubic equation

$$\theta^3 + (x_j - \lambda_n)\theta^2 - (\sigma_1^2 + z_j + x_j \lambda_n)\theta - \sigma_1^2 x_j = 0.$$

We are not able to find here an upper bound on the negative eigenvalues that is strictly smaller than zero. Under our regularity and strict complementarity assumptions, as stated in section 3, $\hat{K}_3$ is nonsingular and converges to a well-defined limit as $\mu \to 0$. Therefore, its condition number is asymptotically uniformly bounded, which is not reflected by the bounds on the eigenvalues in this case. This is where using regularization pays off, since we are able to remove our regularity assumptions and reach

more definitive conclusions regarding the upper bound on the negative eigenvalues and asymptotic condition number estimates.

From Corollary 5.4, the lower bound on negative eigenvalues remains finite asymptotically, while Corollary 5.3 shows that the asymptotic bounds on the positive eigenvalues essentially reduce to those of K_2. In practice, however, we have observed that the condition number of $\hat{K}_3$ is typically substantially better than that of K_1 or K_2 and slightly, but consistently, better than that of K_3.

6. Numerical experiments. We offer a few examples to validate the analysis of the previous sections using a basic MATLAB implementation of Mehrotra's predictor-corrector procedure with an initial point computed as proposed by Friedlander and Orban (2012). The linear systems are solved using MATLAB's backslash operator which performs a factorization. Eigenvalues, singular values, and condition numbers are computed with MATLAB's built-in functions. Our code is not intended to rival state-of-the-art solvers, but simply to illustrate the analysis in previous sections.

We illustrate the eigenvalue bounds for the regularized systems of section 4 and the traditional systems considered in section 5. We use two small convex QPs from the TOMLAB[1] optimization software package, problems number 6 and 21, both converted to standard form. Problem 6 has 10 variables, 7 constraints, is strictly convex, and its solution satisfies strict complementarity and the LICQ. Problem 21 has 51 variables and 27 constraints, and the Hessian is positive semidefinite with rank 3. The solution identified by our interior-point method satisfies strict complementarity but not the LICQ. Those two problems are very small but are nonetheless representative of the situation for several other problems we have tested. The bounds for the regularized systems are illustrated using $\rho = \delta = 10^{-4}$ throughout the interior-point iterations.

We begin with Problem 21. Figure 6.1 shows the eigenvalues of $K_{1,\text{reg}}$ and K_1 on a semilogarithmic scale. The bounds for $K_{1,\text{reg}}$ become nearly fixed after a few initial iterations. On the other hand, for K_1, we note that because H is singular, the large eigenvalues grow without bound and the small eigenvalues converge to zero, in accordance with the bounds of Corollary 5.1.

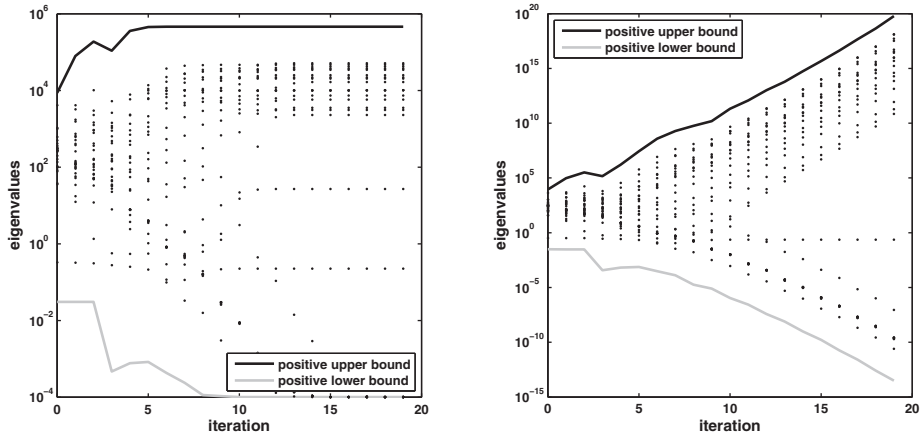

(a) $\rho = \delta = 10^{-4}$ (b) $\rho = \delta = 0$

FIG. 6.1. TOMLAB *problem* 21: *eigenvalues of* $K_{1,reg}$ *and* K_1 *and their bounds.*

[1]Available at tomopt.com/tomlab.

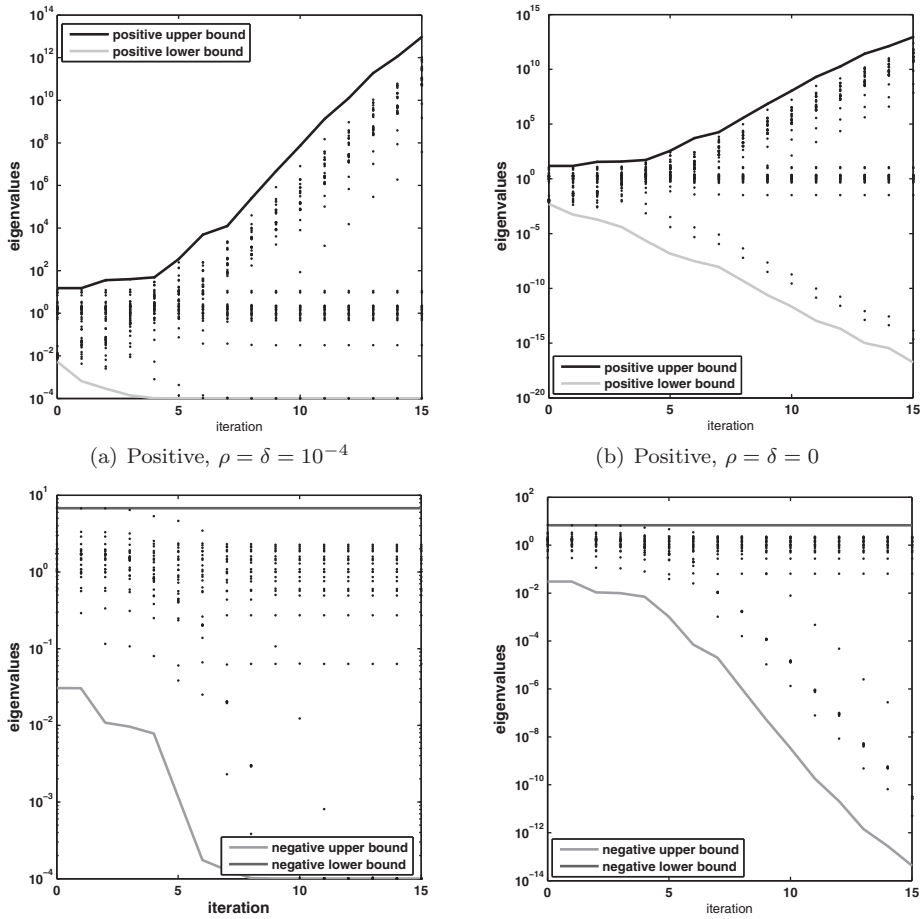

(a) Positive, $\rho = \delta = 10^{-4}$

(b) Positive, $\rho = \delta = 0$

(c) Negative (absolute values), $\rho = \delta = 10^{-4}$

(d) Negative (absolute values), $\rho = \delta = 0$

FIG. 6.2. TOMLAB *problem* 21: *eigenvalues of* $K_{2,reg}$ *and* K_2 *and their bounds.*

Figure 6.2 shows the positive and negative eigenvalues of $K_{2,\mathrm{reg}}$ and K_2 on separate semilogarithmic plots. Note that the absolute value of the negative eigenvalues is represented. As evident from the formulas in Theorem 4.2 and Corollary 5.2, as we approach the solution and x_i or z_i become smaller due to the complementarity condition, the formulas for the bounds on the negative eigenvalues are increasingly prone to cancelation errors. To avoid harmful cancellation errors, we use the mathematically equivalent but more numerically appropriate formulas (which apply to K_2 as well with $\delta = \rho = 0$)

$$\gamma^- \geq -\frac{2\left(\sigma_1^2 + \delta\lambda_{\min}(H_\rho)\right)}{\lambda_{\min}(H_\rho) - \delta + \sqrt{(\lambda_{\min}(H_\rho) - \delta)^2 + 4(\sigma_1^2 + \delta\lambda_{\min}(H_\rho))}},$$

$$\gamma^- \leq -\frac{2\left(\sigma_m^2 + \delta\lambda_{\max}(H_\rho)\right)}{\lambda_{\max}(H_\rho) - \delta + \sqrt{(\lambda_{\max}(H_\rho) - \delta)^2 + 4(\sigma_m^2 + \delta\lambda_{\max}(H_\rho))}}.$$

For $K_{2,\mathrm{reg}}$, while the "inner" bounds (namely, the lower positive and upper negative bounds) are well away from zero, the upper bound on the positive eigenvalues

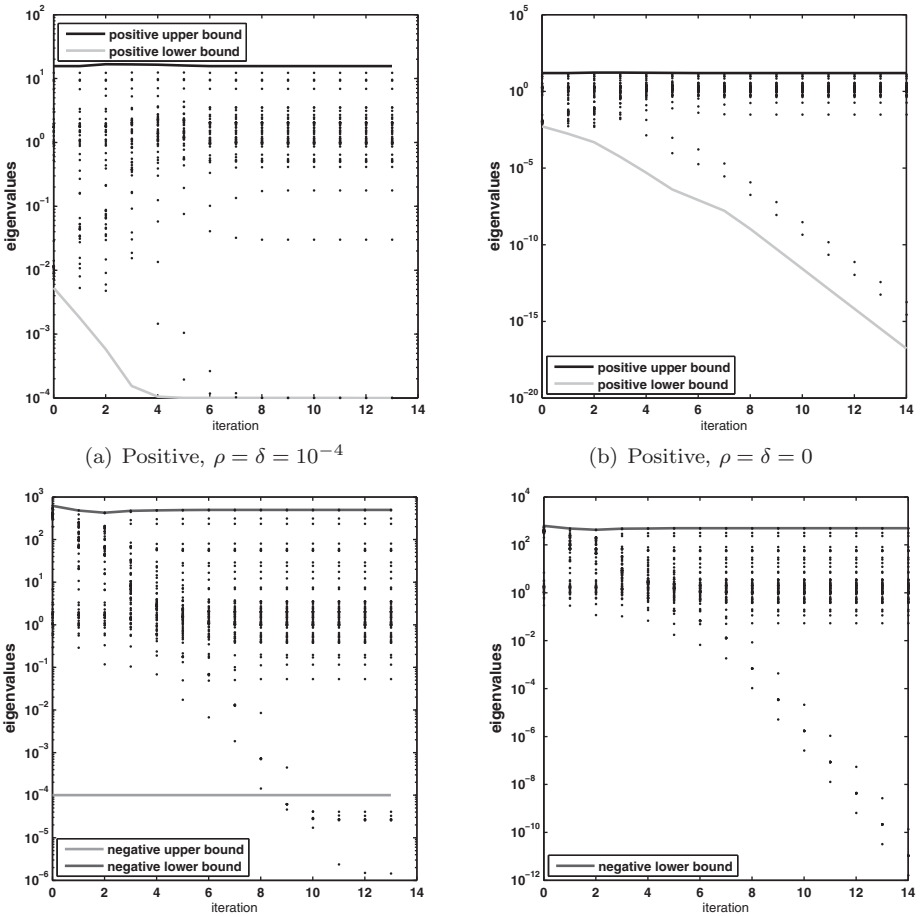

(a) Positive, $\rho = \delta = 10^{-4}$

(b) Positive, $\rho = \delta = 0$

(c) Negative (absolute values), $\rho = \delta = 10^{-4}$

(d) Negative (absolute values), $\rho = \delta = 0$

FIG. 6.3. TOMLAB *problem* 21: *eigenvalues of* $K_{3,reg}$ *and* K_3 *and their bounds. Note that the upper bound on the negative eigenvalues is zero for* K_3 *and is thus not visible in graph* (d).

still increases without bound, as predicted by Theorem 4.2, and all bounds are tight. For K_2, the upper bound on the negative eigenvalues also decays, but the lower bound on the negative eigenvalues is fixed. This is in accordance with Corollary 5.2 and we see that the bounds are relatively tight in this example.

Next, we show results for the 3×3 block system. We choose to show results for the nonsymmetric versions; the eigenvalues for the symmetrized versions were indistinguishable in our computations. Figure 6.3 shows the eigenvalues of $K_{3,\mathrm{reg}}$ and K_3 on separate semilogarithmic plots. For $K_{3,\mathrm{reg}}$, the lower bound on the positive eigenvalues is well away from zero despite the smallest eigenvalue of the Hessian being zero, and the bound seems to converge fairly quickly to $\rho = 10^{-4}$ as iterations progress; this is in line with the results in Theorem 4.3. For K_3, since $\rho = 0$ and we have $\lambda_n = 0$ in this example, the lower bound on the positive eigenvalues in Corollary 5.3 is prone to cancellation errors, since this bound simplifies to $\gamma^+ \geq \min_j \frac{1}{2}(-x_j + \sqrt{x_j^2 + 4z_j})$. To avoid cancellation, in our numerical experiments we use the mathematically equivalent

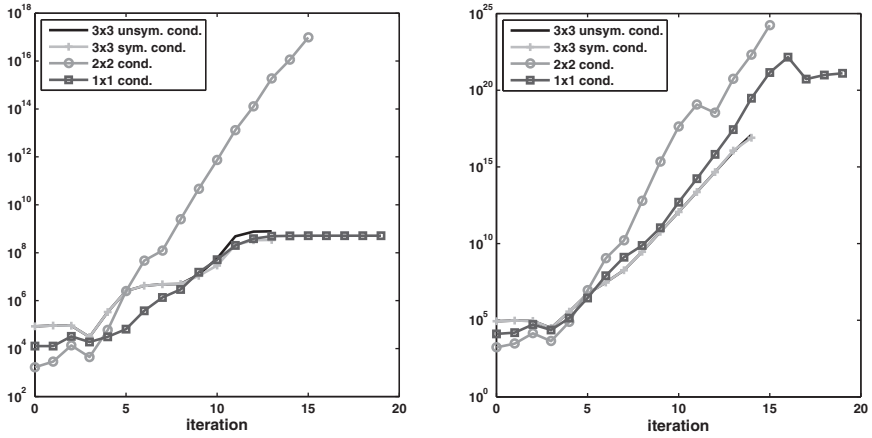

(a) $\rho = \delta = 10^{-4}$ (b) $\rho = \delta = 0$

FIG. 6.4. TOMLAB *problem* 21: *condition numbers for regularized and unregularized systems.*

formula

$$\gamma^+ \geq \min_j \frac{2z_j}{x_j + \sqrt{x_j^2 + 4z_j}}$$

with the convention that the bound is zero if strict complementarity fails to hold.

For the negative eigenvalues, we have an upper bound at $-\delta$ as per Theorem 4.4. Interestingly, at iteration 9 the magnitude of the smallest negative eigenvalues drops below δ, illustrating the "gray zone" when the conditions of Theorem 4.4 no longer hold but Theorem 4.5 does not yet apply. In fact, the assumptions of Theorem 4.5 are not satisfied for this problem. Despite these eigenvalues dropping below the bound, they remain away from zero. For K_3, the two "outer" bounds (that is, the upper positive and lower negative bounds) are fixed and tight. On the other hand, the upper bound on the negative eigenvalues is zero and the lower bound on the positive eigenvalues decays toward zero. The eigenvalues also decay toward zero from both sides. This is in accordance with Corollaries 5.3 and 5.4.

Next, we consider the condition numbers for all formulations and examine the regularized problem with $\rho = \delta = 10^{-4}$ and the unregularized problem. Figure 6.4 shows the condition numbers of the different formulations. With regularization, $K_{1,\text{reg}}$ and both the nonsymmetric and the symmetrized 3×3 formulations have bounded condition numbers, while $K_{2,\text{reg}}$ is still numerically singular. For the unregularized problem, all formulations are numerically singular at the end of the iterations, though K_3 and $\hat{K}_3$ have the best condition number—their curves are nearly superposed. The singularity in this case is due to the fact that this problem does not satisfy the LICQ.

We now consider the condition numbers for all formulations applied to Problem 6. Figure 6.5 shows the condition numbers. The condition number of $K_{1,\text{reg}}$ remains essentially fixed after the initial iterations because of the use of fixed regularization parameters, in accordance with Theorem 4.1. Note that in theory, this condition number could be as large as 10^8 but it stabilizes around 10^5 in this instance. The condition number of $K_{2,\text{reg}}$ goes unbounded with the iterations, as predicted by the asymptotic analysis. For the unregularized problem, K_3 and $\hat{K}_3$ are well-conditioned throughout, while the reduced forms have exponentially increasing condition numbers.

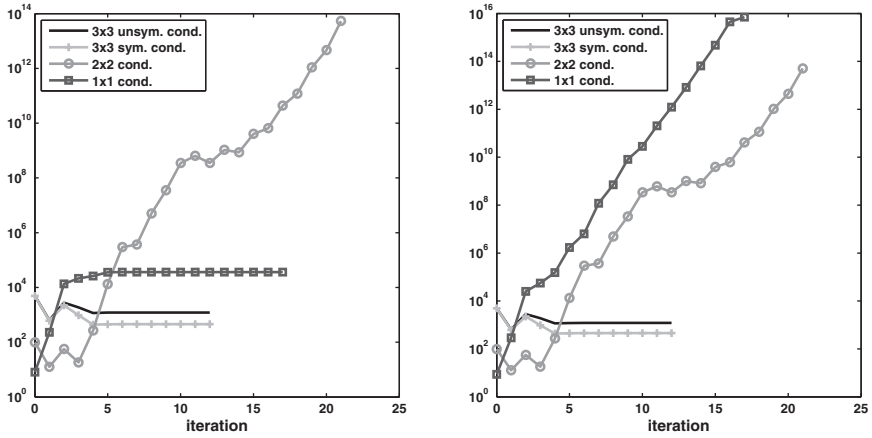

(a) $\rho = \delta = 10^{-4}$ (b) $\rho = \delta = 0$

FIG. 6.5. TOMLAB *problem* 6: *condition numbers for regularized and unregularized systems.*

Since strict complementarity and LICQ are satisfied for this problem, regularization has little effect on the 3×3 formulations.

7. Other formulations. In this section, we briefly mention alternative linear systems that were not covered in the previous sections but to which our techniques appear to generalize. In particular, we examine a formulation adapted to problems that do not satisfy strict complementarity at a solution. We refer below to the unregularized formulation, (1.1a)–(1.1b).

When $H = 0$, the Goldman and Tucker (1956) theorem guarantees that a strictly complementary solution always exists provided there exists at least one solution. Moreover, widely applicable interior-point frameworks guarantee that any limit point of the sequence of iterates determines a strictly complementary solution under mild assumptions—see, e.g., (Wright, 1997, Theorem 6.8). When $H \neq 0$, this desirable result no longer holds in general, i.e., not all problems of the form (1.1a) possess a solution satisfying strict complementarity. A typical counterexample is

$$\underset{x \in \mathbb{R}}{\text{minimize}} \ \tfrac{1}{2}x^2 \quad \text{subject to } \ x \geq 0$$

for which it is easy to verify that the only primal-dual solution is $(x, z) = (0, 0)$. A difficulty with such problems is that, in the limit, K_3 is singular while K_2 and K_1 do not even appear to be well defined. Our test scenario of Example 2.1 no longer describes the general situation because there is a subset of $\mathcal{A}$ for which $x_i = \Theta(\sqrt{\mu})$ and $z_i = \Theta(\sqrt{\mu})$—see, e.g., (Coulibaly, Gould, and Orban, 2012; Monteiro and Wright, 1994; Wright and Orban, 2002).

It is common to study the iterates generated by an interior-point method in the vicinity of a strictly complementary solution. Consider the typical situation where (1.1a) possesses a solution with at least one zero variable to which a positive multiplier is associated. An immediate difficulty is that K_2 does not converge to a well-defined limit and appears to become arbitrarily ill-conditioned—an observation that is confirmed by the results of section 5. The same holds for K_1. We now outline a strategy to salvage the situation in the case of K_2, and that may be applied to K_3 as well. In the vicinity of a strictly complementary solution of (1.1)—assuming one

exists—partition the variables according to $\mathcal{A}$ and $\mathcal{I}$ and consider the induced partitioning of the matrices H, J, X, and Z. The system (2.2) may be written as

$$\begin{bmatrix} H_{\mathcal{A}\mathcal{A}} & H_{\mathcal{A}\mathcal{I}}^T & J_{\mathcal{A}}^T & -I & \\ H_{\mathcal{A}\mathcal{I}} & H_{\mathcal{I}\mathcal{I}} & J_{\mathcal{I}}^T & & -I \\ J_{\mathcal{A}} & J_{\mathcal{I}} & & & \\ -Z_{\mathcal{A}\mathcal{A}} & & & -X_{\mathcal{A}\mathcal{A}} & \\ & -Z_{\mathcal{I}\mathcal{I}} & & & -X_{\mathcal{I}\mathcal{I}} \end{bmatrix} \begin{bmatrix} \Delta x_{\mathcal{A}} \\ \Delta x_{\mathcal{I}} \\ \Delta y \\ \Delta z_{\mathcal{A}} \\ \Delta z_{\mathcal{I}} \end{bmatrix} = \begin{bmatrix} r_{c,\mathcal{A}} \\ r_{c,\mathcal{I}} \\ r_b \\ r_{\tau,\mathcal{A}} \\ r_{\tau,\mathcal{I}} \end{bmatrix}.$$

Note that $X_{\mathcal{I}\mathcal{I}}^{-1} Z_{\mathcal{I}\mathcal{I}}$ approaches zero. Gould (1986) then proposes to eliminate $\Delta z_{\mathcal{I}}$ and reformulate the above system as

$$\begin{bmatrix} H_{\mathcal{A}\mathcal{A}} & H_{\mathcal{A}\mathcal{I}}^T & J_{\mathcal{A}}^T & -I \\ H_{\mathcal{A}\mathcal{I}} & H_{\mathcal{I}\mathcal{I}} + X_{\mathcal{I}\mathcal{I}}^{-1} Z_{\mathcal{I}\mathcal{I}} & J_{\mathcal{I}}^T & \\ J_{\mathcal{A}} & J_{\mathcal{I}} & & \\ -Z_{\mathcal{A}\mathcal{A}} & & & -X_{\mathcal{A}\mathcal{A}} \end{bmatrix} \begin{bmatrix} \Delta x_{\mathcal{A}} \\ \Delta x_{\mathcal{I}} \\ \Delta y \\ \Delta z_{\mathcal{A}} \end{bmatrix} = \begin{bmatrix} r_{c,\mathcal{A}} \\ r_{c,\mathcal{I}} - X_{\mathcal{I}\mathcal{I}}^{-1} r_{\tau,\mathcal{I}} \\ r_b \\ r_{\tau,\mathcal{A}} \end{bmatrix}.$$

In addition, Gould (1986) symmetrizes the matrix by multiplying the last block row by $Z_{\mathcal{A}\mathcal{A}}^{-1}$. Since $X_{\mathcal{A}\mathcal{A}} Z_{\mathcal{A}\mathcal{A}}^{-1}$ also approaches zero, the resulting matrix then possesses a well-defined limit as long as J has full row rank, and therefore we should expect its condition number to be uniformly bounded, at least in a neighborhood of a strictly complementary solution.

This partial elimination requires accurate identification of the index sets $\mathcal{A}$ and $\mathcal{I}$. In practice, this can be done by using the predictive index set $\{i \mid x_i/z_i < z_i/x_i\}$ as an approximation to $\mathcal{A}$.

Some of our results rely on strict complementarity being satisfied in the limit. We now outline a similar partitioning by which our results will apply to problems failing to satisfy strict complementarity.

A possible approach with such problems is to use indicator sets to distinguish between indices $i = 1, \ldots, n$ that are weakly active, strongly active, and inactive. Let (x, y, z) be a local approximation to a solution. The set of *strongly active* constraints at x is $\mathcal{A}_S := \{i = 1, \ldots, n \mid x_i = 0 < z_i\}$. The set of *weakly active* constraints at x is $\mathcal{A}_W := \{i = 1, \ldots, n \mid x_i = z_i = 0\}$ and the set of *inactive* constraints at x is $\mathcal{I} := \{i = 1, \ldots, n \mid z_i = 0 < x_i\}$. Suppose at each iteration k of an interior-point method, we have a mechanism to identify approximations $\mathcal{A}_S^k$, $\mathcal{A}_W^k$, and $\mathcal{I}^k$ to $\mathcal{A}_S$, $\mathcal{A}_W$, and $\mathcal{I}$, respectively. Such indicator sets can resolve the singular limit difficulty provided they ensure that $z_i^k/x_i^k \to 0$ as $k \to \infty$ for $i \in \mathcal{A}_W^k \cup \mathcal{I}^k$ while $x_i^k/z_i^k \to 0$ as $k \to \infty$ for $i \in \mathcal{A}_S^k$. Indeed if this were the case, upon partitioning x, z, H, and J according to $\mathcal{B}^k := \mathcal{A}_W^k \cup \mathcal{I}^k$ and $\mathcal{S}^k := \mathcal{A}_S^k$, (2.2) could be partially eliminated to

$$\begin{bmatrix} H_{\mathcal{S}\mathcal{S}} & H_{\mathcal{S}\mathcal{B}}^T & J_{\mathcal{S}}^T & -I \\ H_{\mathcal{S}\mathcal{B}} & H_{\mathcal{B}\mathcal{B}} + X_{\mathcal{B}\mathcal{B}}^{-1} Z_{\mathcal{B}\mathcal{B}} & J_{\mathcal{B}}^T & \\ J_{\mathcal{S}} & J_{\mathcal{B}} & & \\ -Z_{\mathcal{S}\mathcal{S}} & & & -X_{\mathcal{S}\mathcal{S}} \end{bmatrix} \begin{bmatrix} \Delta x_{\mathcal{S}} \\ \Delta x_{\mathcal{B}} \\ \Delta y \\ \Delta z_{\mathcal{S}} \end{bmatrix} = \begin{bmatrix} r_{c,\mathcal{S}} \\ r_{c,\mathcal{B}} + X_{\mathcal{B}\mathcal{B}}^{-1} r_{\tau,\mathcal{B}} \\ r_b \\ r_{\tau,\mathcal{S}} \end{bmatrix}.$$

The matrix of the latter system has a well-defined limit whenever J has full row rank. As above, the typical way to symmetrize the system is to multiply the last block row

by $Z_{\mathcal{SS}}^{-1}$. Examples of indicator sets with the requisite properties along with pointers to the literature are given by Coulibaly, Gould, and Orban (2012) and Monteiro and Wright (1994).

In both cases of partial elimination above, should J not have full row rank, dual regularization, as described in section 2, will ensure that the partially eliminated matrix converges to a well-defined limit, provided the dual regularization parameter remains bounded away from zero.

We also note that our results of section 3 can be extended to the appropriate symmetrization of the two above partially eliminated systems with the benefit of solving smaller systems with similar properties.

There are other formulations of the system that we have not covered in this paper. Forsgren and Gill (1998) propose a *doubly augmented* formulation that has the benefit of being positive definite on and near the central path, allowing the use of more specialized linear solvers. Korzak (1999) and Fourer and Mehrotra (1993) both use scalings of the systems to alleviate ill-conditioning. Gill et al. (1992) proceed similarly to Gould (1986) and partition inequality constraints into active and inactive constraints, and note that ill-conditioning is due to the varying sizes of constraint values across the partition. Benzi, Haber, and Taralli (2010) discuss the conditioning of the formulations that we cover, focusing on a specific application. They use iterative methods with large-scale problems, where the ill-conditioning of the reduced forms is problematic. In their experiments, they use the original unreduced form and a partially eliminated formulation.

8. Concluding remarks. Our analysis indicates that 3×3 block linear systems that arise throughout the iterations of primal-dual interior-point solvers have a favorable spectral structure. The matrix is nonsingular, whereas it tends to singularity when it is reduced using commonly used partial elimination procedures. Regularization is shown to be very effective at shifting the eigenvalues and the eigenvalue bounds away from zero.

Our "outer" bounds, namely, the upper bound on the positive eigenvalues and the lower bound on the negative eigenvalues, seem tight. On the other hand, the "inner" bounds, namely, the lower bound on the positive eigenvalues and the upper bound on the negative eigenvalues, are pessimistic in some cases. In the regularized case we have found all bounds, and consequently were able to establish an asymptotic condition number. Numerical experiments confirm that the largest eigenvalues in absolute value of the unreduced 3×3 block matrix are bounded away from zero and only modestly grow throughout the iterations. Analysis for the upper negative bound is challenging, and our results involve a region that we term the "gray zone," which represents the transition from the course of iterations to the limit. Our analysis does not provide an effective bound for the gray zone, but we believe that by tightening our assumptions we will be able to further shrink this zone.

By continuity of the eigenvalues, the analysis for the unregularized systems follows as a special case of the regularized systems, simply by setting the regularization parameters to zero. Here we have been able to find three effective bounds, but our bound on the negative eigenvalues is zero, even though the matrix is analytically known to be nonsingular.

It is important to note that convergence may be driven by several considerations other than the eigenvalue distribution, especially when direct linear solvers are used. The subspace where the right-hand side lies and the scaling of the problem are two of several factors that affect convergence.

We have not addressed the issue of the performance of iterative solvers, but our work in this paper has been done with our eyes on this paradigm. As problems become larger in scale, the importance of studying these solvers increases. For such solvers the distribution of eigenvalues plays a central role, and the need to find effective preconditioners is critical and challenging. The question of how to utilize inexactness throughout the iterations is also important. Benzi, Haber, and Taralli (2010), Bergamaschi et al. (2007), and Forsgren, Gill, and Griffin (2007) provide relevant and important observations and results. The bounds given in this paper as well as preliminary experimentation lead us to believe that the 3×3 formulation has a strong potential to be very effective in the context of iterative solvers.

Matrices similar to K_1, K_2, and K_3 occur in other branches of optimization, such as cone optimization. The difference in this case is that X and Z may no longer be diagonal but are instead restricted to being symmetric and positive definite. In addition, they may be dense. Whether our results generalize to such a scenario may be an interesting subject for future research.

Acknowledgments. We are very grateful to the two referees for their insightful comments, which have significantly improved the contents and the structure of this paper. We also wish to thank Jacek Gondzio and Michael Saunders for carefully reading this paper and for numerous valuable comments. The third author gratefully acknowledges constant moral support from Ludwig Van Beethoven.

REFERENCES

A. Altman and J. Gondzio (1999), *Regularized symmetric indefinite systems in interior point methods for linear and quadratic optimization*, Optim. Methods Softw., 11, pp. 275–302.

P. Armand and J. Benoist (2011), *Uniform boundedness of the inverse of a Jacobian matrix arising in regularized interior-point methods*, Math. Program., 137, pp. 587–592.

O. Axelsson and M. Neytcheva (2006), *Eigenvalue estimates for preconditioned saddle point matrices*, Numer. Linear Algebra Appl., 13, pp. 339–360.

Z.-Z. Bai (2013), *Eigenvalue estimates for saddle point matrices of Hermitian and indefinite leading blocks*, J. Comput. Appl. Math., 237, pp. 295–306.

Z.-Z. Bai, M. K. Ng, and Z.-Q. Wang (2009), *Constraint preconditioners for symmetric indefinite matrices*, SIAM J. Matrix Anal. Appl., 31, pp. 410–433.

M. Benzi, G. H. Golub, and J. Liesen (2005), *Numerical solution of saddle point problems*, Acta Numer., 14, pp. 1–137.

M. Benzi, E. Haber, and L. Taralli (2010), *Multilevel algorithms for large-scale interior point methods*, SIAM J. Sci. Comput., 31, pp. 4152–4175.

L. Bergamaschi, J. Gondzio, M. Venturin, and G. Zilli (2007), *Inexact constraint preconditioners for linear systems arising in interior point methods*, Comput. Optim. Appl., 36, pp. 137–147.

R. H. Byrd, M. E. Hribar, and J. Nocedal (1999), *An interior point algorithm for large-scale nonlinear programming*, SIAM J. Optim., 9, pp. 877–900.

R. H. Byrd, J.-Ch. Gilbert, and J. Nocedal (2000), *A trust region method based on interior point techniques for nonlinear programming*, Math. Program. Ser. A, 89, pp. 149–185.

Z. Coulibaly, N. I. M. Gould, and D. Orban (2012), *Fast local convergence of interior-point methods in the absence of strict complementarity*, manuscript, in preparation.

J. Czyzyk, S. Mehrotra, M. Wagner, and S. J. Wright (1999), *PCx: An interior-point code for linear programming*, Optim. Methods Softw., 11, pp. 397–430.

A. V. Fiacco and G. P. McCormick (1990), *Nonlinear programming*, 2nd ed., Classics Appl. Math. 4, SIAM, Philadelphia.

A. Forsgren (2002), *Inertia-controlling factorizations for optimization algorithms*, Appl. Numer. Math., 43, pp. 91–107.

A. Forsgren and P. E. Gill (1998), *Primal-dual interior methods for nonconvex nonlinear programming*, SIAM J. Optim., 8, pp. 1132–1152.

A. Forsgren, P. E. Gill, and J. R. Shinnerl (1996), *Stability of symmetric ill-conditioned systems arising in interior methods for constrained optimization*, SIAM J. Matrix Anal. Appl., 17, pp. 187–211.

A. Forsgren, P. E. Gill, and M. H. Wright (2002), *Interior methods for nonlinear optimization*, SIAM Rev., 44, pp. 525–597.

A. Forsgren, P. E. Gill, and J. D. Griffin (2007), *Iterative solution of augmented systems arising in interior methods*, SIAM J. Optim., 18, pp. 666–690.

R. Fourer and S. Mehrotra (1993), *Solving symmetric indefinite systems in an interior-point method for linear programming*, Math. Program., 62, pp. 15–39.

M. P. Friedlander and D. Orban (2012), *A primal-dual regularized interior-point method for convex quadratic programs*, Math. Program. Comput., 4, pp. 71–107.

E. M. Gertz and S. J. Wright (2003), *Object-oriented software for quadratic programming*, ACM Trans. Math. Software, 29, pp. 58–81.

P. E. Gill, W. Murray, D. B. Ponceleón, and M. A. Saunders (1992), *Preconditioners for indefinite systems arising in optimization*, SIAM J. Matrix Anal. Appl., 13, pp. 292–311.

A. J. Goldman and A. W. Tucker (1956), *Theory of linear programming*, in Linear Equalities and Related Systems, H. W. Kuhn and A. W. Tucker, eds., Princeton University Press, Princeton, NJ, pp. 53–97.

J. Gondzio (2012), *Matrix-free interior point methods*, Comput. Optim. Appl., 51, pp. 457–480.

N. I. M. Gould (1985), *On practical conditions for the existence and uniqueness of solutions to the general equality quadratic-programming problem*, Math. Program., 32, pp. 90–99.

N. I. M. Gould (1986), *On the accurate determination of search directions for simple differentiable penalty functions*, IMA J. Numer. Anal., 6, pp. 357–372.

N. I. M. Gould and V. Simoncini (2010), *Spectral analysis of saddle point matrices with indefinite leading blocks*, SIAM J. Matrix Anal. Appl., 31, pp. 1152–1171.

J. Korzak (1999), *Eigenvalue relations and conditions of matrices arising in linear programming*, Computing, 62, pp. 45–54.

S. Mehrotra (1992), *On the implementation of a primal-dual interior point method*, SIAM J. Optim., 2, pp. 575–601.

R. D. C. Monteiro and S. J. Wright (1994), *Local convergence of interior-point algorithms for degenerate monotone LCP*, Comput. Optim. Appl., 3, pp. 131–155.

E. Moulding (2012), *Spectral Properties of Matrices Arising from Primal-Dual Interior-Point Methods for Convex Quadratic Programs*, Master's thesis, The University of British Columbia, Vancouver, Canada.

N. M. Nachtigal, S. C. Reddy, and L. N. Trefethen (1992), *How fast are nonsymmetric matrix iterations?*, SIAM J. Matrix Anal. Appl., 13, pp. 778–795.

T. Rusten and R. Winther (1992), *A preconditioned iterative method for saddlepoint problems*, SIAM J. Matrix Anal. Appl., 13, pp. 887–904.

Y. Saad (2003), *Iterative Methods for Sparse Linear Systems*, SIAM, Philadelphia.

M. A. Saunders (1996), *Cholesky-based methods for sparse least squares: The benefits of regularization*, in Linear and Nonlinear Conjugate Gradient-Related Methods, L. Adams and J. L. Nazareth, eds., SIAM, Philadelphia, pp. 92–100.

D. Silvester and A. Wathen (1994), *Fast iterative solution of stabilised Stokes systems part* II. *Using general block preconditioners*, SIAM J. Numer. Anal., 31, pp. 1352–1367.

R. J. Vanderbei (1995), *Symmetric quasidefinite matrices*, SIAM J. Optim., 5, pp. 100–113.

A. Wächter and L. T. Biegler (2006), *On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming*, Math. Program. Ser. A, 106, pp. 25–57.

J. H. Wilkinson (1959), *The evaluation of the zeros of ill-conditioned polynomials. Part* I, Numer. Math., 1, pp. 150–166.

M. H. Wright (1992), *Interior methods for constrained optimization*, Acta Numer., 1, pp. 341–407.

M. H. Wright (1994), *Some properties of the Hessian of the logarithmic barrier function*, Math. Program., 67, pp. 265–295.

M. H. Wright (2005), *The interior-point revolution in optimization: History, recent developments, and lasting consequences*, Bull. Amer. Math. Soc., 42, pp. 39–56.

S. J. Wright (1997), *Primal-Dual Interior-Point Methods*, SIAM, Philadelphia.

S. J. Wright and D. Orban (2002), *Local convergence of the Newton/log-barrier method for degenerate problems*, Math. Oper. Res., 27, pp. 585–613.