

Inexact primal-dual path-following algorithms for a special class of convex quadratic SDP and related problems

K.C. Toh ^{*}, R. H. Tütüncü [†] and M. J. Todd [‡]

December 30, 2005

Dedicated to Masakazu Kojima on the occasion of his 60th birthday

Abstract

We propose a primal-dual path-following Mehrotra-type predictor-corrector method for solving convex quadratic semidefinite programming (QSDP) problems. For the special case when the quadratic term has the form $\frac{1}{2}X \bullet (UXU)$, we compute the search direction at each iteration from the Schur complement equation. We are able to solve the Schur complement equation efficiently via the preconditioned symmetric quasi-minimal residual iterative solver with two appropriately constructed preconditioners. Numerical experiments on a variety of QSDPs with matrices of dimensions up to 2000 are performed and the computational results show that our methods are efficient and robust. Our methods can also be extended to linear SDP problems with upper bound constraints on primal matrix variables.

Keywords: semidefinite programming, semidefinite least squares, path-following methods, Nesterov-Todd scaling, symmetric quasi-minimum residual iteration

AMS subject classification: 90C22, 90C25, 90C51, 65F10

^{*}Department of Mathematics, National University of Singapore, 2 Science Drive 2, Singapore 117543 (mattohk@nus.edu.sg).

[†]Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA (reha@cmu.edu). Research supported in part by NSF through grants CCR-9875559, CCF-0430868 and by ONR through grant N00014-05-1-0147

[‡]School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York 14853, USA (miketodd@cs.cornell.edu). Research supported in part by NSF through grant DMS-0209457 and by ONR through grant N00014-02-1-0057.

1 Introduction

Let \mathcal{S}^n denote the space of real symmetric matrices of order n . Our interest is in interior-point methods for convex quadratic semidefinite programming problems, i.e., optimization problems over \mathcal{S}^n in which the matrix variable is required to be positive semidefinite, where now the usual linear objective function is augmented by a quadratic function of the symmetric matrix variable. Specifically, we consider the following convex quadratic semidefinite program (QSDP)

$$\begin{aligned} (QSDP) \quad & \min_X \frac{1}{2}X \bullet \mathcal{Q}(X) + C \bullet X \\ & \mathcal{A}(X) = b, \quad X \succeq 0, \end{aligned} \tag{1}$$

where $\mathcal{Q} : \mathcal{S}^n \rightarrow \mathcal{S}^n$ is a given self-adjoint positive semidefinite operator on \mathcal{S}^n and $\mathcal{A} : \mathcal{S}^n \rightarrow \mathbb{R}^m$ is a linear map. The notation $X \succeq 0$ indicates that X is in \mathcal{S}_+^n , the cone of positive semidefinite real symmetric matrices of order n , and $U \bullet V$ denotes $\text{Tr}(U^T V)$, the usual trace inner product. The adjoint of \mathcal{A} with respect to the standard inner products in \mathcal{S}^n and \mathbb{R}^m is denoted by \mathcal{A}^T . The dual of (1) is given as follows:

$$\begin{aligned} (QSDD) \quad & \max_{X,y,S} -\frac{1}{2}X \bullet \mathcal{Q}(X) + b^T y \\ & \mathcal{A}^T(y) - \mathcal{Q}(X) + S = C, \quad S \succeq 0. \end{aligned} \tag{2}$$

QSDPs first appeared in the literature in the work of Kojima, Shindoh, and Hara [21] as a special case of monotone semidefinite linear complementarity problems (SDLCPs). Kojima and his co-authors showed that SDLCPs can be reduced to standard SDPs [22]. Since this reduction is often inefficient (see Section 2), a separate study of SDLCPs and QSDPs is justified. In Section 2, we describe some applications leading to problems of the form (QSDP) and previous work on algorithms for such problems.

Another form of quadratic SDP has been considered in [7], namely, $\min_y \{ \frac{1}{2}y^T H y + b^T y : \mathcal{A}^T y \preceq C, y \in \mathbb{R}^m \}$, where H is a given positive semidefinite matrix. In this case, the Schur complement matrix arising at each interior-point iteration has the form $H + \mathcal{A}\mathcal{E}^{-1}\mathcal{F}\mathcal{A}^T$ (with \mathcal{E} and \mathcal{F} as in Section 3), and the computation presents no difficulties, being very similar to that for a standard linear SDP. As our interest is in problems with quadratic terms involving the matrix variables, we shall not consider this form of quadratic SDP further.

In this paper, we propose a primal-dual path-following Mehrotra-type predictor-corrector method for (1) and (2). For a general self-adjoint positive semidefinite \mathcal{Q} , the search direction at each iteration must be computed from an augmented system of dimension $m + n(n+1)/2$. Our ultimate goal is to investigate the efficient computation of the search direction by applying a preconditioned symmetric quasi-minimal residual (PSQMR) method to the augmented system. We focus our attention on the efficient computation of the search direction only for the important special case where \mathcal{Q} has the form $\mathcal{Q} = U \otimes U$, with U being a given matrix in \mathcal{S}_+^n . We show that the search direction can be computed from the much smaller Schur complement system of dimension m in this case. The cost is comparable to that of computing the AHO direction [3] for a standard linear SDP. We also discuss the conditioning of this Schur complement matrix asymptotically.

As the cost of computing the search direction in the special case can still be very expensive, we propose to use a PSQMR method to solve the governing Schur complement equation. When we use Nesterov-Todd scaling to compute the search directions, the Schur complement matrix has the form $M = \mathcal{A}(G_1 \otimes G_1 + G_2 \otimes G_2)^{-1} \mathcal{A}^T$, where G_1, G_2 are symmetric positive definite matrices. We show that $(G_1 \otimes G_1 + G_2 \otimes G_2)^{-1}$ admits a semi-analytical expression of the form $\mathcal{J} \mathcal{D} \mathcal{J}^T$. We propose two preconditioners for M . The first preconditioner is based on a low rank approximation of the diagonal operator (with respect to the usual basis) \mathcal{D} . The second preconditioner is constructed by approximating the sum of two Kronecker products by a single Kronecker product.

It turns out that a Schur complement matrix of the form given in the last paragraph also arises at each interior-point iteration for a linear SDP with an upper bound. Thus we also apply our PSQMR method to solve such a problem.

Our interest here is in the efficient computation of the search direction at each iteration and we do not explicitly address the issue of the computational complexity of our algorithm. However, it is worth pointing out that for the more general monotone semidefinite complementarity problem, Kojima, Shindoh, and Hara [21] showed that their path-following method using the dual-HKM direction has polynomial complexity. A similar result would hold for the HKM direction. Our algorithm differs in that we use the NT direction (which has advantages in computing the search direction) and also Mehrotra's predictor-corrector approach.

The rest of this article is organized as follows. As already mentioned, Section 2 focuses on examples of QSDPs and earlier algorithmic studies of these problems. In Section 3, we describe our primal-dual interior-point algorithm. Sections 4 and 5 contain our main contributions; namely the detailed discussion of the efficient computation of the search direction for general \mathcal{Q} and structured \mathcal{Q} , respectively, including the preconditioners used in the solution of large Schur complement systems. Section 6 discusses the upper-bounded linear SDP and its relation to QSDPs. Finally, in Section 7, we provide numerical results for a variety of QSDP problems discussed in Sections 2 and 6, with matrices of order up to 2000. The results are encouraging.

We employ the following notation in this paper. Given an integer n , we let $\bar{n} = n(n+1)/2$. Consider the isometry $\mathbf{svec} : \mathcal{S}^n \rightarrow \mathbb{R}^{\bar{n}}$ defined by

$$\mathbf{svec}(X) = [X_{11}, \sqrt{2}X_{12}, X_{22}, \dots, \sqrt{2}X_{1n}, \dots, \sqrt{2}X_{n-1,n}, X_{nn}].$$

We write the matrix representation of $\mathcal{Q}(X)$ in the standard basis of \mathcal{S}^n as $\mathbf{svec}(\mathcal{Q}(X)) = Q\mathbf{svec}(X)$, where Q is a positive semidefinite matrix in $\mathcal{S}^{\bar{n}}$; similarly, the matrix representation of $\mathcal{A}(X)$ is written as $\mathcal{A}(X) = A\mathbf{svec}(X)$, where A is a matrix in $\mathbb{R}^{m \times \bar{n}}$. We use $\|\cdot\|_2$ to denote either the vector 2-norm or the matrix 2-norm. The Frobenius norm of a matrix is denoted by $\|\cdot\|_F$. Given any square matrices P and Q of order n , $P \otimes Q$ denotes the symmetrized Kronecker product operator on \mathcal{S}^n defined by $P \otimes Q(M) = (QMP^T + PMQ^T)/2$. For details on the properties of \otimes and its relation to the standard Kronecker product, see the Appendix of [36].

2 Examples and Existing Work

We start by noting that QSDPs can be reformulated and solved as standard (i.e., linear objective) semidefinite-quadratic-linear programming (SQLP) problems [37]. Consider the Cholesky factorization $Q = R^T R$, where $R \in \mathbb{R}^{p \times \bar{n}}$, with p being the rank of Q . (Note that when Q has full rank, $p = \bar{n}$.) It is readily shown that (1) can be reformulated as an SQLP after introducing $p + 1$ artificial variables and p linear constraints as follows:

$$\min \left\{ \frac{1}{2}t + C \bullet X : \begin{bmatrix} A \\ R \end{bmatrix} \text{svec}(X) + \begin{bmatrix} 0 & 0 \\ 0 & -I \end{bmatrix} \begin{bmatrix} t \\ s \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}, X \succeq 0, \|s\|_2^2 \leq t \right\}, \quad (3)$$

where the constraint $\|s\|_2^2 \leq t$ can easily be converted into a standard second order cone constraint. The computational cost required to solve the reformulated problem (3) grows at least like $O((m+p)^3)$ and the memory requirement grows like $O((m+p)^2)$. Thus, unless $m + p$ is small it is extremely expensive to solve (QSDP) by reformulating it into a standard SQLP. Given that p can be as large as $\bar{n} = \Theta(n^2)$, it is safe to say that an approach based on (3) can comfortably solve only problems with n at most 100 on a high end PC available today. Therefore, while QSDPs are equivalent to a simpler class of problems, the reformulation is costly and inefficient, and a study dedicated to alternative algorithmic approaches for QSDPs is justified.

For general QSDPs we have already mentioned the work of Kojima and his co-authors in [21, 22]. These authors made additional contributions to the theoretical study of algorithms for SDLCPs in follow-up articles [23, 24]. Another study that addresses general QSDPs is given in [32]. These authors consider an interior-point algorithm based on reducing a primal-dual potential function. Their algorithm has an iteration complexity of $O(\sqrt{\bar{n}} \ln(1/\epsilon))$ for computing an ϵ -optimal solution. At each iteration, the search direction needs to be computed from an augmented system of dimension $m + \bar{n}$. As the linear system is generally very large, the authors proposed using the conjugate gradient (CG) method to compute an approximate direction, but no preconditioning was discussed although it is crucial to do so to ensure that the CG method has a reasonable convergence rate. Furthermore, the authors do not report any numerical implementation to test the performance of their proposed method.

One of the most common types of QSDPs encountered in the literature is the linearly constrained semidefinite least squares problem:

$$(SDLS) \quad \min_X \left\{ \left\| \mathcal{L}(X) - \widehat{K} \right\|_F : \mathcal{A}(X) = b, X \succeq 0 \right\}, \quad (4)$$

where $\mathcal{L} : \mathcal{S}^n \rightarrow \mathcal{S}^p$ is a linear map and \widehat{K} is a given symmetric matrix in \mathcal{S}^p . Very often, $\widehat{K} = \mathcal{L}(K)$ for some $K \in \mathcal{S}^n$, and then the objective function is $\|\mathcal{L}(X - K)\|_F$: we seek a positive semidefinite X satisfying certain linear restrictions and as close as possible to K in some suitable weighted sense. However, the EDM problem below uses a more general \widehat{K} as above.

In finding the nearest correlation matrix to a given data matrix K [18], $\mathcal{A}(X) = b$ represents the constraints that fix the diagonal elements of X to one, and \mathcal{L} (with $p = n$)

typically has the form $\mathcal{L}(X) = U^{1/2}XU^{1/2}$ with a symmetric positive definite weight matrix U , or $\mathcal{L}(X) = U \circ X$, with a symmetric elementwise positive weight matrix U . Here and below, the notation “ \circ ” means element-wise matrix multiplication. In the QSDP formulation of these problems, \mathcal{Q} takes the form $U \circledast U$ in the first case and the element-wise multiplication operator with the matrix $U \circ U$ in the second case. It is worth noting that the operator in the second case is positive definite and diagonal. For these problems, Higham proposes and analyzes a modified alternating projection solution method.

For the special case of the unweighted nearest correlation matrix problem for which $\mathcal{L}(X) = X$ (and the corresponding operator \mathcal{Q} in (1) is the identity), Anjos et al. [6] proposed a feasible primal-dual interior-exterior algorithm for (1) based on inexact Gauss-Newton directions computed from overdetermined systems each with n^2 equations and \bar{n} unknowns. Preconditioned CG methods with diagonal and block diagonal incomplete Cholesky preconditioning are used to compute the inexact directions. Unfortunately, preliminary numerical results obtained by the inexact Gauss-Newton approach do not show that it is numerically more efficient than the standard formulation (3).

In [28], Malick proposed a partial Lagrangian dual algorithm for solving the nearest correlation matrix problem by dualizing the linear constraints. A quasi-Newton method is used to solve the dual problem $\max\{b^T y - \|\mathcal{P}(\mathcal{A}^T y - C)\|_F^2 : y \in \mathbb{R}^m\}$, where $\mathcal{P}(U)$ is the projection of U onto \mathcal{S}_+^n . According to the numerical results reported in [28], this method performs very well on the nearest correlation matrix problem. More recently, Boyd and Xiao [8], apparently unaware of the work in [28], also proposed a Lagrangian dual approach combined with a projected sub-gradient method to solve the nearest correlation matrix problem.

In [33], Qi and Sun proposed a non-smooth Newton method for the same problem. Based on recent results on strongly semismooth matrix valued functions, they were able to establish quadratic convergence of their method. Numerical experiments in [33] show that the non-smooth Newton method is highly effective.

SDLS also arises from the problem of finding the nearest Euclidean distance matrix for a given weighted graph $G = (V, E, \omega)$ on p nodes [4]. Let \widehat{K}_G be a $p \times p$ matrix whose elements are given by $k_{ij} = \omega_{ij}$ if $(i, j) \in E$, and $k_{ij} = 0$ if $(i, j) \notin E$. We seek points q_1, q_2, \dots, q_p in \mathbb{R}^n for $n = p - 1$ such that $\|q_i - q_j\|^2$ is close to k_{ij} for $(i, j) \in E$. Then the optimization problem (see [4]) is the following:

$$(EDM) \quad \min_X \left\{ \left\| \mathcal{L}(X) - \widehat{K}_G \right\|_F : \mathcal{A}(X) = b, X \succeq 0 \right\}, \quad (5)$$

where now $\mathcal{L}(X) = \Sigma \circ (\text{diag}(VXV^T)e^T + e \text{diag}(VXV^T)^T - 2VXV^T)$. Here Σ is the adjacency matrix of the graph G , and V is a $p \times n$ matrix such that $V^T e = 0$ and $V^T V = I_n$. If we factor the solution X as RR^T , then the q 's can be taken to be the columns of $R^T V^T$.

For semidefinite least squares problems (4) arising from the nearest Euclidean distance matrix problem Alfakih et al. [4] proposed a primal-dual interior-point algorithm based on the Gauss-Newton approach to solve the perturbed optimality conditions. Unfortunately, the linear system that needs to be solved at each iteration has dimension

about n^2 . Consequently, only small problems with $n \leq 50$ can be comfortably solved on a standard PC.

3 A primal-dual path-following interior-point algorithm

We propose to solve the primal-dual pair (1) and (2) using primal-dual path-following methods based on their perturbed KKT conditions:

$$\begin{aligned} -\mathcal{Q}(X) + \mathcal{A}^T(y) + S &= C, \quad S \succeq 0 \\ \mathcal{A}(X) &= b, \quad X \succeq 0 \\ XS &= \nu I, \end{aligned} \tag{6}$$

where $\nu > 0$ is a positive parameter. Given the current iterate (X, y, S) with X and S positive definite, the search direction $(\Delta X, \Delta y, \Delta S)$ at each interior-point iteration is the solution of the following symmetrized Newton system:

$$\begin{aligned} -\mathcal{Q}(\Delta X) + \mathcal{A}^T(\Delta y) + \Delta S &= R_d := C - S - \mathcal{A}^T y + \mathcal{Q}(X) \\ \mathcal{A}(\Delta X) &= r_p := b - \mathcal{A}(X) \\ \mathcal{E}(\Delta X) + \mathcal{F}(\Delta S) &= R_c := \sigma \mu I - H_P(XS), \end{aligned} \tag{7}$$

where \mathcal{E} and \mathcal{F} are linear operators on S^n that depend on the symmetrization scheme $H_P(\cdot)$ chosen, with P being the symmetrization matrix; for more details, see for example [31], [36]. Here, $\mu = X \bullet S/n$, and $\sigma \in (0, 1)$ is the centering parameter.

We start with an initial iterate (X^0, y^0, S^0) with $X^0, S^0 \succ 0$, and step-length parameter $\tau^0 = 0.9$. The details of an iteration of our primal-dual path-following algorithm are explained below. In this description we denote the current and the next iterates by (X, y, S) and (X^+, y^+, S^+) , and the current and the next step-length parameters by τ and τ^+ , respectively.

Algorithm IP-QSDP.

- Set $\mu = X \bullet S/n$.
- (Convergence test)

Stop the iteration if the accuracy measure ϕ is sufficiently small, where

$$\phi = \max \left\{ \frac{X \bullet S}{1 + |\text{pobj}| + |\text{dobj}|}, \frac{\|r_p\|_2}{1 + \|b\|_2}, \frac{\|R_d\|_F}{1 + \|C\|_F} \right\} \tag{8}$$

with r_p, R_d defined as in (7), $\text{pobj} = \frac{1}{2}X \bullet \mathcal{Q}(X) + C \bullet X$, and $\text{dobj} = -\frac{1}{2}X \bullet \mathcal{Q}(X) + b^T y$.

- (Predictor step)

Compute the predictor search direction $(\delta X, \delta y, \delta S)$ from (7) by choosing $\sigma = 0$.

- (Predictor step-length)
Compute

$$\alpha_p = \min(1, \tau \alpha). \quad (9)$$

Here α is the maximum step length that can be taken so that $X + \alpha\delta X$ and $S + \alpha\delta S$ remain positive semidefinite.

- (Centering rule)
Set $\sigma = (X + \alpha_p\delta X) \bullet (S + \alpha_p\delta S) / X \bullet S$.
- (Corrector step)
Compute the search direction $(\Delta X, \Delta y, \Delta S)$ from (7), with R_c replaced by

$$R'_c = \sigma\mu I - H_P(XS + \delta X\delta S).$$

- (Corrector step-length)
Compute α_c as in (9) but with $(\delta X, \delta S)$ replaced by $(\Delta X, \Delta S)$.
- Update (X, y, S) to (X^+, y^+, S^+) by

$$X^+ = X + \alpha_c \Delta X, \quad y^+ = y + \alpha_c \Delta y, \quad S^+ = S + \alpha_c \Delta S. \quad (10)$$

- Update the step-length parameter by $\tau^+ = 0.9 + 0.08\alpha_c$.

4 Computation of search directions: general \mathcal{Q}

The linear system of equations (7) is a non-symmetric system of dimension $m+n(n+1)$. This is generally a very large system even for a moderate n , say, $n = 200$. Thus it is extremely expensive to solve the system directly. By eliminating ΔS , we get the following augmented equation with dimension $m + \bar{n}$:

$$\underbrace{\begin{bmatrix} -\mathcal{H} & A^T \\ \mathcal{A} & 0 \end{bmatrix}}_{\mathcal{B}} \begin{bmatrix} \Delta X \\ \Delta y \end{bmatrix} = \begin{bmatrix} R_d - \mathcal{F}^{-1}R_c \\ r_p \end{bmatrix}, \quad (11)$$

where

$$\mathcal{H} = \mathcal{F}^{-1}\mathcal{E} + \mathcal{Q}. \quad (12)$$

Note that under the assumptions that \mathcal{A} is a surjective map and $X, S \succ 0$, the matrix \mathcal{B} is nonsingular when the scaling matrix P is chosen to be in the Monteiro-Zhang family described in [31]. Specific members of the family include the Nesterov-Todd scaling [36] and the HKM scaling, independently proposed by Helmberg et al. [17] and Kojima et al. [21], and later rederived from a different viewpoint by Monteiro [30]. For the NT scaling, we have $\mathcal{F}^{-1}\mathcal{E} = W^{-1} \otimes W^{-1}$, where W is the unique matrix in \mathcal{S}_+^n satisfying $WSW = X$. For the HKM scaling, we have $\mathcal{E}^{-1}\mathcal{F} = X \otimes S^{-1}$, but

unfortunately, $\mathcal{F}^{-1}\mathcal{E}$ does not have such a simple analytical expression. For the dual HKM scaling, we do not have a simple expression for $\mathcal{E}^{-1}\mathcal{F}$, but $\mathcal{F}^{-1}\mathcal{E} = S \circledast X^{-1}$.

By further eliminating ΔX , we get the Schur complement equation of dimension m below:

$$\underbrace{\mathcal{A}\mathcal{H}^{-1}\mathcal{A}^T}_M \Delta y = h := r_p + \mathcal{A}\mathcal{H}^{-1}(R_d - \mathcal{F}^{-1}R_c). \quad (13)$$

For a general \mathcal{Q} , even for the simple case where \mathcal{Q} is a diagonal operator, \mathcal{H} cannot be inverted at moderate cost, so the computation of the Schur complement matrix M is extremely expensive. Thus, unlike the case of linear SDP, computing the direction based on (13) is computationally not feasible with the possible exception of the case when \mathcal{Q} is of the form $U \circledast U$. The best alternative seems to be to compute the direction based on the augmented equation (11). Using (11) instead of (13), we avoid the costly construction of the matrix M . However, the coefficient matrix \mathcal{B} in (11) is typically very large and the matrix \mathcal{H} in its (1,1) block is typically dense. As a result, solving (11) by a direct method such as the LDL^T factorization method is out of consideration. It is necessary to use an iterative solver such as the preconditioned symmetric quasi-minimal residual (PSQMR) method [11] to solve (11).

We will discuss in the next section the special case where $\mathcal{Q} = U \circledast U$ for which \mathcal{H}^{-1} has an analytical expression, and the computational cost of the Schur complement matrix is more moderate.

Remark 4.1 *To avoid the need to handle the sum $\mathcal{F}^{-1}\mathcal{E} + \mathcal{Q}$ whose inverse is expensive to compute, we explore another route starting from the augmented system (11). Consider the Cholesky factorization $\mathcal{Q} = \mathcal{R}^T\mathcal{R}$, where the matrix representation of \mathcal{R} is a matrix in $\mathbb{R}^{p \times \bar{n}}$. Let $\Delta Z = -\mathcal{R}\Delta X$. Then the augmented system can be rewritten as follows:*

$$\begin{bmatrix} -\mathcal{F}^{-1}\mathcal{E} & \mathcal{A}^T & \mathcal{R}^T \\ \mathcal{A} & 0 & 0 \\ \mathcal{R} & 0 & I \end{bmatrix} \begin{bmatrix} \Delta X \\ \Delta y \\ \Delta Z \end{bmatrix} = \begin{bmatrix} R_d - \mathcal{F}^{-1}R_c \\ r_p \\ 0 \end{bmatrix}. \quad (14)$$

The introduction of the auxiliary variable ΔZ is motivated by the paper [26] for convex quadratic programming in \mathbb{R}^n . Upon eliminating ΔX from (14), we get the following linear system of dimension $m + p$:

$$\begin{bmatrix} M_{AA} & M_{RA}^T \\ M_{RA} & I + M_{RR} \end{bmatrix} \begin{bmatrix} \Delta y \\ \Delta Z \end{bmatrix} = \begin{bmatrix} h_y \\ h_Z \end{bmatrix} \quad (15)$$

where $M_{AA} = \mathcal{A}\mathcal{E}^{-1}\mathcal{F}\mathcal{A}^T$, $M_{RA} = \mathcal{R}\mathcal{E}^{-1}\mathcal{F}\mathcal{A}^T$, $M_{RR} = \mathcal{R}\mathcal{E}^{-1}\mathcal{F}\mathcal{R}^T$, and

$$h_y = r_p + \mathcal{A}\mathcal{E}^{-1}\mathcal{F}R_d - \mathcal{A}\mathcal{E}^{-1}R_c, \quad h_Z = \mathcal{R}\mathcal{E}^{-1}\mathcal{F}R_d - \mathcal{R}\mathcal{E}^{-1}R_c.$$

Unfortunately, the system (15) has exactly the same dimension and structure as the potentially very large Schur complement matrix arising at each interior-point iteration for the standard SQLP (3). In particular, note that the matrices M_{AA}, M_{RA}, M_{RR} are

typically all dense, even if \mathcal{R} is a diagonal operator. Thus, unless $m + p$ is small, the approach based on (15) would not offer any computational advantage over the augmented system (11). As such, we shall not consider this approach any further in this paper.

When using an iterative method to solve (11), it is necessary to know how accurately one must solve the equation. For this purpose, we will construct a stopping condition based on the residual norm of the computed search direction with respect to (7).

Proposition 4.1 (a) Suppose $(\Delta X, \Delta y)$ is computed from (11) with residual vector given by (η_1, η_2) . Suppose that ΔS is computed exactly from

$$\Delta S = R_d - \mathcal{A}^T \Delta y + \mathcal{Q}(\Delta X).$$

Then the residual vector associated with the direction $(\Delta X, \Delta y, \Delta S)$ for (7) is given by $(0, \eta_2, -\mathcal{F}(\eta_1))$.

(b) Suppose Δy is computed from (13) with residual vector given by ξ . Suppose that ΔX and ΔS are computed exactly as follows:

$$\Delta X = \mathcal{H}^{-1} \left(\mathcal{A}^T \Delta y - (R_d - \mathcal{F}^{-1} R_c) \right), \quad \Delta S = R_d - \mathcal{A}^T \Delta y + \mathcal{Q}(\Delta X). \quad (16)$$

Then the residual vector associated with the direction $(\Delta X, \Delta y, \Delta S)$ for (7) is given by $(0, \xi, 0)$.

Proof. We omit the proof since it is straightforward. □

We deem $(\Delta X, \Delta y)$ computed from (11) to be sufficiently accurate if the following relative stopping condition is satisfied:

$$\|(\eta_2, -\mathcal{F}(\eta_1))\|_2 \leq \kappa \max\{\|r_p\|_2, \|R_d\|_F, \|R_c\|_F\},$$

where $\kappa \in (0, 1)$ is an accuracy parameter. Similarly, we deem Δy computed from (13) to be sufficiently accurate if $\|\xi\|_2 \leq \kappa \max\{\|r_p\|_2, \|R_d\|_F, \|R_c\|_F\}$. In the numerical experiments in Section 7, we choose $\kappa = 10^{-3}$.

4.1 Conditioning of M

The convergence rate of a Krylov subspace method such as the SQMR method depends heavily on the condition number of the coefficient matrix of the linear system being solved. Thus it is of interest to analyze the conditioning of the coefficient matrix M in (13). For simplicity, we assume in this subsection that the NT scaling is used, that is,

$$\mathcal{H} = W^{-1} \circledast W^{-1} + \mathcal{Q}. \quad (17)$$

This assumption is without loss of generality: for points on the central path, which we consider below, the NT scaling, HKM scaling, and dual HKM scaling coincide. We will

show that, under suitable additional assumptions, the condition number of M remains bounded on the central path of (1) and (2).

Assume that the problems (1) and (2) are strictly feasible, and that \mathcal{A} is surjective. These are necessary and sufficient conditions for the existence and uniqueness of solutions (X_ν, y_ν, S_ν) of the central path equations (6). Also, these solutions converge to some optimal solution (X_*, y_*, S_*) as ν tends to zero; see Halicka, de Klerk, and Roos [14], and Luo, Sturm, and Zhang [27]. Both papers additionally contain many results on the behavior of X_ν and S_ν . We further assume that X_* is primal nondegenerate and that strict complementarity holds in the sense of Alizadeh, Haeberly, and Overton [2] (see below). Thus (y_*, S_*) is the unique optimal dual solution, and the ranks of X_* and S_* sum to n . We will show that, under these conditions, the condition number of M_ν , the Schur complement matrix in (13) corresponding to (X_ν, y_ν, S_ν) , remains bounded as ν tends to zero.

Let us suppose not, and choose a monotonically decreasing sequence $\{\nu_k\}$ such that $\lim_{k \rightarrow \infty} \nu_k = 0$ and $\lim_{k \rightarrow \infty} \|M_{\nu_k}\|_2 \|M_{\nu_k}^{-1}\|_2 = \infty$. For simplicity of notation, we write M_k, X_k, S_k , etc., for $M_{\nu_k}, X_{\nu_k}, S_{\nu_k}$, and so on. Since X_k and S_k commute, there is an orthogonal matrix P_k that simultaneously diagonalizes X_k and S_k so that

$$X_k = P_k \Lambda_k P_k^T, \quad S_k = P_k \Sigma_k P_k^T,$$

where the eigenvalue matrices

$$\Lambda_k = \text{Diag}(\lambda_1^k, \dots, \lambda_n^k), \quad \Sigma_k = \text{Diag}(\sigma_1^k, \dots, \sigma_n^k)$$

satisfy $\lambda_i^k \sigma_i^k = \nu_k$, and the eigenvalues are ordered such that

$$\lambda_1^k \geq \dots \geq \lambda_n^k > 0, \quad 0 < \sigma_1^k \leq \dots \leq \sigma_n^k.$$

Let P_* be a limit point of the set $\{P_k\}$. We refine the sequence if necessary so that $\{P_k\}$ converges to P_* . Then P_* is an orthogonal matrix that simultaneously diagonalizes X_* and S_* with

$$X_* = P_* \Lambda_* P_*^T, \quad S_* = P_* \Sigma_* P_*^T, \tag{18}$$

where

$$\Lambda_* = \text{Diag}(\lambda_1^*, \dots, \lambda_n^*), \quad \Sigma_* = \text{Diag}(\sigma_1^*, \dots, \sigma_n^*)$$

satisfy $\lambda_i^* \sigma_i^* = 0$, and

$$\lambda_1^* \geq \dots \geq \lambda_r^* > \lambda_{r+1}^* = \dots = \lambda_n^* = 0, \quad 0 = \sigma_1^* = \dots = \sigma_{n-s}^* < \sigma_{n-s+1}^* \leq \dots \leq \sigma_n^*,$$

where r and s are the ranks of X_* and S_* , respectively. We are assuming that (X^*, y^*, S^*) satisfies the strict complementarity condition, i.e., that $r + s = n$. Let $P_{*,1}$ and $P_{*,2}$ be the submatrices denoting the first r and the last $n - r$ columns of P_* , respectively. Suppose the symmetric matrices A_k are defined by $\mathcal{A}^T y = \sum_{k=1}^m y_k A_k$. We are also assuming that X^* is primal nondegenerate [2], i.e., that the set

$$\left\{ [P_{*,1}^T A_k P_{*,1}, P_{*,1}^T A_k P_{*,2}; P_{*,2}^T A_k P_{*,1}, 0] : k = 1, \dots, m \right\}$$

is linearly independent in \mathcal{S}^n (this is an equivalent definition: see Theorem 6 in [2]).

By the strict complementarity assumption, $r + s = n$. For k sufficiently large, the NT scaling matrix associated with (X_k, S_k) , given by $W_k = P_k D_k^{-1} P_k^T$ with $D_k = \Lambda_k^{-1/2} \Sigma_k^{1/2}$, has r and s eigenvalues of the orders $\Theta(1/\sqrt{\nu_k})$ and $\Theta(\sqrt{\nu_k})$, respectively. This implies that $W_k^{-1} \circledast W_k^{-1}$ has \bar{r} , rs , and \bar{s} eigenvalues of the orders $\Theta(\nu_k)$, $\Theta(1)$, and $\Theta(1/\nu_k)$, respectively.

Recall that M_k denotes the Schur complement matrix in (13) corresponding to (X_k, y_k, S_k) . We observe that

$$M_k = \mathcal{A} \mathcal{P}_k \left(D_k \circledast D_k + \tilde{\mathcal{Q}}_k \right)^{-1} \mathcal{P}_k^T \mathcal{A}^T,$$

where $\mathcal{P}_k = P_k \circledast P_k$, and $\tilde{\mathcal{Q}}^k = \mathcal{P}_k^T \mathcal{Q} \mathcal{P}_k$. Assume that the eigenvalues of $W_k^{-1} \circledast W_k^{-1}$ are sorted in increasing order. Consider the following partition

$$D_k \circledast D_k + \tilde{\mathcal{Q}}_k = \begin{bmatrix} \mathcal{D}_1^k + \tilde{\mathcal{Q}}_{11}^k & \tilde{\mathcal{Q}}_{12}^k \\ \tilde{\mathcal{Q}}_{21}^k & \tilde{\mathcal{Q}}_{22}^k + \mathcal{D}_2^k \end{bmatrix}, \quad (19)$$

where \mathcal{D}_1^k and \mathcal{D}_2^k correspond to diagonal matrices whose diagonal elements consist of the first $\bar{r} + rs$, and the last \bar{s} eigenvalues of $W_k^{-1} \circledast W_k^{-1}$, respectively. Note that there is a positive constant ρ such that $0 \prec \mathcal{D}_1^k \preceq \rho I$ for all $k = 1, \dots$, and that the diagonal entries of the matrix corresponding to \mathcal{D}_2^k are all $\Theta(1/\nu_k)$.

Let $\mathcal{P}_{k,1}$ be the submatrix obtained from \mathcal{P}_k by extracting its first $\bar{r} + rs$ columns. Similarly, let $\mathcal{P}_{*,1}$ be the submatrix obtained from $\mathcal{P}_* := P_* \circledast P_*$ by extracting its first $\bar{r} + rs$ columns, and $\tilde{\mathcal{Q}}^* = \mathcal{P}_*^T \mathcal{Q} \mathcal{P}_*$, with $\tilde{\mathcal{Q}}_{11}^*$ its part corresponding to $\tilde{\mathcal{Q}}_{11}^k$. We have the following theorem.

Theorem 4.1 *Suppose the problems (1) and (2) are strictly feasible, and that strict complementarity holds for the optimal solution (X_*, y_*, S_*) . Suppose further that $\tilde{\mathcal{Q}}_{11}^*$ is positive definite. Let $M_* = \mathcal{A} \mathcal{P}_{*,1} (\tilde{\mathcal{Q}}_{11}^*)^{-1} \mathcal{P}_{*,1}^T \mathcal{A}^T$. Then the following results hold:*

- (a) *There is a positive constant c_1 such that $c_1 \limsup_{k \rightarrow \infty} \|M_k\|_2 \leq \|M_*\|_2$.*
- (b) *Suppose in addition that primal nondegeneracy holds for the optimal solution (X_*, y_*, S_*) . Then there exists a positive constant c_2 such that $\limsup_{k \rightarrow \infty} \|M_k^{-1}\|_2 \leq c_2 \|M_*^{-1}\|_2 < \infty$.*

Proof. Since $\lim_{k \rightarrow \infty} \tilde{\mathcal{Q}}_{11}^k = \tilde{\mathcal{Q}}_{11}^*$ (which is positive definite) and $0 \prec \mathcal{D}_1^k \preceq \rho I$ for all $k = 1, \dots$, there exist constants $c_1, c_2 > 0$ such that for sufficiently large k , we have $c_1 \tilde{\mathcal{Q}}_{11}^* \preceq \mathcal{D}_1^k + \tilde{\mathcal{Q}}_{11}^k \preceq c_2 \tilde{\mathcal{Q}}_{11}^*$. This implies that $c_1 (\mathcal{D}_1^k + \tilde{\mathcal{Q}}_{11}^k)^{-1} \preceq (\tilde{\mathcal{Q}}_{11}^*)^{-1} \preceq c_2 (\mathcal{D}_1^k + \tilde{\mathcal{Q}}_{11}^k)^{-1}$.

By using the formula for the inverse of a 2×2 block matrix in [34, p.389], one can show from (19) that for sufficiently large k ,

$$\left(D_k \circledast D_k + \tilde{\mathcal{Q}}_k \right)^{-1} = \begin{bmatrix} (\mathcal{D}_1^k + \tilde{\mathcal{Q}}_{11}^k)^{-1} & 0 \\ 0 & 0 \end{bmatrix} + O(\nu_k). \quad (20)$$

Thus for sufficiently large k ,

$$M_k = \mathcal{A}\mathcal{P}_{k,1} \left(\mathcal{D}_1^k + \tilde{\mathcal{Q}}_{11}^k \right)^{-1} \mathcal{P}_{k,1}^T \mathcal{A}^T + O(\nu_k),$$

and

$$c_1(M_k + O(\nu_k)) \preceq \mathcal{A}\mathcal{P}_{k,1}(\tilde{\mathcal{Q}}_{11}^*)^{-1}\mathcal{P}_{k,1}^T\mathcal{A}^T \preceq c_2(M_k + O(\nu_k)). \quad (21)$$

Now the result in (a) follows readily from the left-hand side partial order in (21). To prove (b), we show that the middle matrix in (21) is positive definite for sufficiently large k , which follows if M_* is. But this holds as long as $\mathcal{A}\mathcal{P}_{*,1}$ is surjective. Suppose the symmetric matrices A_i are defined by $\mathcal{A}^T y = \sum_i y_i A_i$. The surjectivity condition is equivalent to saying that the matrices $P_*^T A_i P_*$, when their trailing $s \times s$ submatrices are replaced by zeroes (this operation corresponds to considering only the first $\bar{r} + rs$ columns of P_*), are linearly independent. But this follows from the primal nondegeneracy of X_* . Now, by considering the smallest eigenvalues of the matrices on the right-hand side partial order of (21) and taking the limit infimum, we get $\lambda_{\min}(M_*) \leq c_2 \liminf_{k \rightarrow \infty} \lambda_{\min}(M_k)$. The required result follows by noting that the 2-norm of the inverse of a symmetric positive definite matrix is equal to the reciprocal of its smallest eigenvalue. \square

Corollary 4.1 *Under the assumptions of the theorem, the condition number of M_ν remains bounded as ν tends to zero.*

Proof. The result follows straightforwardly from Theorem 4.1. Indeed, we have shown that the assumption that $\|M_\nu\|_2 \|M_\nu^{-1}\|_2$ is unbounded leads to a contradiction. \square

Motivated by a result in the paper [33], we can show that the primal nondegeneracy assumption used in Theorem 4.1(b) holds when \mathcal{A} is the diagonal map and $b > 0$.

Proposition 4.2 *Consider the linear map $\mathcal{A} : \mathcal{S}^n \rightarrow \mathbb{R}^n$ defined by $\mathcal{A}(X) = \text{diag}(X)$, and assume that $b > 0$. Then X^* is primal nondegenerate for any optimal solution (X^*, y^*, S^*) of (1) and (2).*

Proof. First observe that $X = \text{diag}(b)$ is strictly primal feasible. Also, $X = 0$, $y = (\lambda_{\min}(C) - 1)e$, and $S = C - \text{diag}(y)$ is strictly dual feasible, where e is a vector of ones and λ_{\min} denotes the smallest eigenvalue. Therefore, (1) and (2) are both strictly feasible and must have optimal solutions.

We note that the adjoint $\mathcal{A}^T : \mathbb{R}^n \rightarrow \mathcal{S}^n$ is given by $\mathcal{A}^T(y) = \text{Diag}(y)$. Suppose P_* and Λ_* are defined as in (18). Since $b = \mathcal{A}(X_*) = \text{diag}(P_* \Lambda_* P_*^T)$, we have $b = \text{diag}(\sum_{i=1}^r \lambda_i^* P_{*,i} P_{*,i}^T)$, where $P_{*,i}$ denotes the i th column of P_* and again r is the rank of X_* .

Consider the linear system $(P_*^T \text{Diag}(y) P_*)_{ij} = 0$ for all $1 \leq i \leq r$, $1 \leq j \leq n$; or equivalently, $P_{*,i}^T \text{Diag}(y) P_* = 0$ for all $1 \leq i \leq r$. We want to show that y must be zero, so that by the results of [2] X_* is primal nondegenerate. We have $P_{*,i}^T \text{Diag}(y) = 0$ for all $1 \leq i \leq r$ since P_* is nonsingular. Taking the transpose yields $\text{Diag}(y) P_{*,i} = 0$ for all $1 \leq i \leq r$. Multiply this equation on the right by $\lambda_i^* P_{*,i}^T$ and sum over i to get $\text{Diag}(y) (\sum_{i=1}^r \lambda_i^* P_{*,i} P_{*,i}^T) = 0$. Taking the diagonal of this matrix gives $\text{Diag}(y) b = 0$. The assumption made on b implies that we have $y = 0$, and thus X^* is primal nondegenerate. \square

5 Computation of search directions when $\mathcal{Q} = U \circledast U$

For the special case where $\mathcal{Q} = U \circledast U$ with $U \succeq 0$, it is possible to compute the inverse of the mapping \mathcal{H} given in (12) at a more moderate cost if the NT direction is used. The motivation for choosing the NT scaling instead of other scalings in the Monteiro-Zhang family will become clear after we have presented Lemma 5.1 below. Essentially, the matrix \mathcal{H} turns out to be cheaply invertible in this case so that we can use the much smaller Schur complement system (13) instead of (11) to solve for the search directions. Note that for such a choice of \mathcal{Q} and scaling, we have

$$X \bullet \mathcal{Q}(X) = \left\| U^{1/2} X U^{1/2} \right\|_F^2,$$

and \mathcal{H} has the form

$$\mathcal{H} = U \circledast U + W^{-1} \circledast W^{-1}. \quad (22)$$

The inverse of such an \mathcal{H} can be computed via one of the two procedures described in the following lemma.

Lemma 5.1 *Let \mathcal{H} be as in (22). Then*

$$\mathcal{H}^{-1} = (P \circledast P) \left(\mathcal{I} + D \circledast D \right)^{-1} (P \circledast P)^T, \quad (23)$$

where \mathcal{I} is the identity operator on \mathcal{S}^n and P and the diagonal matrix D are computed in one of the two following ways.

- a) Suppose the Cholesky-like factorization $W = R^T R$ and the eigenvalue decomposition $RUR^T = QDQ^T$ are computed. Then set $P = R^T Q$.
- b) Assume that U is positive definite. Suppose the Cholesky factorization $U = L^T L$ and the eigenvalue decomposition $LWL^T = Q\hat{D}Q^T$ are computed. Then set $P = L^{-1}Q$ and $D = \hat{D}^{-1}$.

Proof.

- a) Given any $V \in \mathcal{S}^n$, $\mathcal{H}Y = V$ implies that

$$UYU + W^{-1}YW^{-1} = V.$$

Thus

$$(RUR^T)(R^{-T}YR^{-1})(RUR^T) + R^{-T}YR^{-1} = RVR^T.$$

Let $\tilde{Y} = R^{-T}YR^{-1}$. Then we have

$$(RUR^T)\tilde{Y}(RUR^T) + \tilde{Y} = (RVR^T).$$

With the above identity and the eigenvalue decomposition of RUR^T , it is readily shown that

$$Y = (R^T Q \circledast R^T Q) \left(\mathcal{I} + D \circledast D \right)^{-1} (R^T Q \circledast R^T Q)^T V,$$

and the required result is shown.

b) We first note that $W^{-1} = L^T Q D Q^T L$. \mathcal{H} can be written as follows (see also the multiplication formulas for \otimes provided in the Appendix of [36]):

$$\begin{aligned}\mathcal{H} &= W^{-1} \otimes W^{-1} + U \otimes U = L^T Q D Q^T L \otimes L^T Q D Q^T L + L^T L \otimes L^T L \\ &= (L^T Q \otimes L^T Q) (D \otimes D + I \otimes I) (L^T Q \otimes L^T Q)^T.\end{aligned}$$

Now, since $(G \otimes G)^{-1} = G^{-1} \otimes G^{-1}$ for an invertible G and $P = (L^T Q)^{-T}$, (23) follows easily. □

Since a Cholesky-like factorization, $W = R^T R$, is typically performed in the process of computing W (see [36]), the additional work required for the first method in Lemma 5.1 is only the computation of the matrix RUR^T , its eigenvalue decomposition, and the matrix product $P = R^T Q$. In contrast, for the second method, one needs to compute (but only once) an additional Cholesky factorization of the U matrix, the matrix product LWL^T and its eigenvalue decomposition as in the first method, and $P = L^{-1}Q$, which takes comparable work to the product $R^T Q$ in the first method. Thus the work required in the two methods is comparable. An important exception is the case when U is the identity matrix. Then, the first method requires the computation of the dense matrix products RR^T and $R^T Q$ in addition to the eigenvalue decomposition required for both methods. In any case, the differences between the flop counts required for these two methods will be relatively insignificant given the more expensive parts of the iteration, such as the computation of the M matrix.

Using (23), the Schur complement matrix M in (13) becomes:

$$M := \mathcal{A} \mathcal{H}^{-1} \mathcal{A}^T = \mathcal{A} (P \otimes P) (\mathcal{I} + D \otimes D)^{-1} (P \otimes P)^T \mathcal{A}^T, \quad (24)$$

where the term $\mathcal{I} + D \otimes D$ is a positive definite diagonal operator. The complexity of computing M is $4mn^3 + \frac{1}{2}m^2n^2$ floating point operations if sparsity in \mathcal{A} is totally ignored; see [29]. But even if sparsity in \mathcal{A} is exploited, the structural formula in (24) makes it non-conducive for one to apply the techniques presented in [12] by Fujisawa, Kojima, and Nakata to exploit the sparsity; thus the computational complexity is generally not much lower, and a savings of at most 50% is typical. (The problem is that, if A_i denotes \mathcal{A}^T times a unit vector, it appears necessary to compute *all* entries of $P^T A_i P$ rather than just those corresponding to a nonzero in some A_h , as in linear SDP.) Note that when $m \approx n$, the computational complexity grows like $O(n^4)$.

Remark 5.1 (a) *Our ability to compute \mathcal{H}^{-1} via the semi-analytical formulas presented in Lemma 5.1 depend critically on \mathcal{H} having the form $\mathcal{H} = U \otimes U + G \otimes G$. Even for a slight change to the form $\mathcal{H} = U \otimes U + G \otimes K$, the technique used in the proof of Lemma 5.1 would fail, and we are not aware of the existence of an analogous semi-analytical formula. This is the reason for our focus on the special case $\mathcal{Q} = U \otimes U$ as well as the use of the Nesterov-Todd scaling in the computation of the search directions.*

(b) The matrix of the form given in (24) also arises from the Schur complement equation when solving a standard linear SDP by primal-dual interior point methods using the AHO direction [3]. The Schur complement matrix in that case has the form

$$\mathcal{A}(P \otimes P) \left(I \otimes D \right)^{-1} (P \otimes P)^T \mathcal{B}^T, \quad (25)$$

where $\mathcal{B} = \mathcal{A}(I \otimes X)$, P is orthogonal, and D is diagonal. Thus, the cost of computing M in (24) is comparable to that of (25) for the AHO direction associated with a standard linear SDP.

(c) When the search direction is computed from (24) at each interior-point iteration, the full eigenvalue decompositions of two $n \times n$ dense matrices are required: one for computing the NT scaling matrix and the other for obtaining the semi-analytical formula of \mathcal{H}^{-1} . Each of these eigenvalue decompositions has a computational complexity of $26/3n^3$ [10]. This can easily become the dominant computational cost when n is large because it is difficult to exploit sparsity in computing a full eigenvalue decomposition. For the machine that we use to conduct the numerical experiments in Section 7, the computation of such a decomposition for a dense 2000×2000 matrix takes about a minute.

To illustrate part (b) of Remark 5.1 further we examine the calculation of the search directions using formulas given in Lemma 5.1 in a bit more detail. The similarities with the calculation of the AHO direction will become apparent. For this purpose, we recall the nearest correlation matrix problem we discussed in Section 2. To keep things simple, we do not consider a weighting matrix. Given a $n \times n$ symmetric matrix K , the nearest correlation matrix to K can be found by solving the following problem:

$$\begin{aligned} \min \quad & \frac{1}{2} X \bullet X + C \bullet X \\ \text{s.t.} \quad & E_{ii} \bullet X = 1, \quad i = 1, \dots, n, \\ & X \succeq 0, \end{aligned} \quad (26)$$

where $C = -K$ and E_{ii} is the $n \times n$ matrix whose only non-zero entry is a 1 in the (i, i) position. We now describe the computation of the Schur complement matrix $M = \mathcal{A}\mathcal{H}^{-1}\mathcal{A}^T$ for this problem using (23) and part (b) of the lemma. First, note that $U = I$, so we choose $L = I$. Let $Q\hat{D}Q^T = W$ be the eigenvalue decomposition of W and $D = \hat{D}^{-1}$. Observe that (23) simplifies to:

$$\mathcal{H}^{-1} = (Q \otimes Q) (\mathcal{I} + D \otimes D)^{-1} (Q^T \otimes Q^T).$$

Next, we observe that $(Q^T \otimes Q^T)(E_{ii}) = Q^T E_{ii} Q = Q_{i,:}^T Q_{i,:}$ where $Q_{i,:}$ denotes the i^{th} row of Q . Therefore, the i^{th} column of $M := \mathcal{A}\mathcal{H}^{-1}\mathcal{A}^T$ can be found as follows:

1. $T_i^1 = Q_{i,:}^T Q_{i,:}$.
2. $[T_i^2]_{jk} = \frac{[T_i^1]_{jk}}{1 + d_j d_k}$, $j, k = 1, \dots, n$.
3. $T_i^3 = Q T_i^2 Q^T$.

4. $M_{:,i} = \text{diag}(T_i^3)$.

Above, the d_i 's are the diagonal elements of the matrix D , i.e., the reciprocals of the eigenvalues of the W matrix and the T_i^j 's for $j = 1, 2, 3$ are temporary matrices that can be discarded after Step 4. While the first and second steps of this computation heavily exploited the sparsity structure of \mathcal{A} to achieve $\Theta(n^2)$ floating point operations (flops), the third step however can longer benefit from the sparsity structure and it needs $\Theta(n^3)$ flops. Thus the overall complexity of computing the Schur complement matrix is $\Theta(n^4)$.

Since only the main diagonal elements of T_i^3 are eventually needed in computing $M_{:,i}$, it appears that the $\Theta(n^3)$ complexity in Step 3 can be reduced by computing just the diagonal elements of T_i^3 . However, the alternative approach of computing the diagonal elements via the formula $Q_{j,:} T_i^2 Q_{j,:}^T$, $j = 1, \dots, n$, also needs $\Theta(n^3)$ flops.

5.1 Preconditioners for M , part I

Let $\tilde{\mathcal{A}} = \mathcal{A}(P \otimes P)$, and consider the matrix in (24):

$$M = \tilde{\mathcal{A}}(\mathcal{I} + D \otimes D)^{-1} \tilde{\mathcal{A}}^T. \quad (27)$$

Here P and D are computed as in part (a) or part (b) of Lemma 5.1; in the first case PP^T is easily seen to be W , while in the second it is U^{-1} , so in either case $(P \otimes P)(P \otimes P)^T = PP^T \otimes PP^T$ is easily obtained. The complexity of computing the matrix in (27) is comparable to that for the AHO direction because sparsity in \mathcal{A} cannot be exploited fully. To lower the computational complexity in solving (13), an alternative is to use the PSQMR method with an appropriate preconditioner. In this case, the matrix M need not be constructed explicitly, and only matrix-vector multiplications of the form My for a given y are required. It is easy to show, based on the formula for \mathcal{H}^{-1} given in Lemma 5.1, that each matrix-vector product My costs $mn^2 + 4n^3$ flops if sparsity in \mathcal{A} is totally ignored. We observe that a matrix product of the form PVP^T for a symmetric V can be computed at a cost of $2n^3$ flops; see [29, Remark A.10].

We note that for solving a symmetric positive definite system, the PCG method is most commonly used, but since PSQMR has very similar convergence behavior and computational cost as PCG, we will continue to use PSQMR here.

We try to approximate the middle term $(\mathcal{I} + D \otimes D)^{-1}$ in (27) as the sum of a small number of terms of the form $\Lambda_k \otimes \Lambda_k$. Specifically, for a fixed positive integer q , suppose that Λ_k 's, $k = 1, \dots, q$, are diagonal matrices such that

$$\sum_{k=1}^q \alpha_k \Lambda_k \otimes \Lambda_k \approx (\mathcal{I} + D \otimes D)^{-1}, \quad (28)$$

where each α_k is a scalar. We will make clear the meaning of “ \approx ” later. Let $V_k = P\Lambda_k P^T$. Suppose $G = (P \otimes P)(\sum_{k=1}^q \alpha_k \Lambda_k \otimes \Lambda_k)(P \otimes P)^T = \sum_{k=1}^q \alpha_k V_k \otimes V_k$. Then G is an approximation to \mathcal{H}^{-1} . Thus, it is natural to consider the following preconditioner

for M :

$$\widehat{M} = \sum_{k=1}^q \alpha_k \mathcal{A}(V_k \otimes V_k) \mathcal{A}^T. \quad (29)$$

The complexity for computing \widehat{M} is at most q times that of the Schur complement matrix associated with the NT direction for a standard linear SDP. The cost of computing each term $\mathcal{A}(V_k \otimes V_k) \mathcal{A}^T$ is $2mn^3 + \frac{1}{2}m^2n^2$ if sparsity in \mathcal{A} is ignored [29], [39]. But we should emphasize that in (29), sparsity in \mathcal{A} can fully be exploited using the techniques in [12], which is not the case for the matrix M in (27). Thus, the computational cost of \widehat{M} is potentially much lower than that of M . It is precisely the difference between (29) and (27) in exploiting sparsity that motivated us to consider the preconditioner (29).

For the approximation problem (28), one can consider minimizing the 2-norm of the difference of the two diagonal matrices on both sides of the approximation. We focus on the vectors consisting of their diagonal elements and then consider the symmetric matrices obtained from these vectors as follows. The elements of each vector (of length \bar{n}) are extracted sequentially to fill in the upper-triangular part of an $n \times n$ matrix column-wise, and the lower-triangular part is then filled in to yield a symmetric matrix. For the right-hand-side matrix in (28), this process yields $K := 1./(ee^T + dd^T)$, the matrix whose entries are the reciprocals of the entries of $ee^T + dd^T$, where $d = \text{diag}(D)$ and e is the vector of ones. For the left-hand-side matrix, the corresponding term is $\sum_{k=1}^q \alpha_k \lambda_k \lambda_k^T$ where $\lambda_k = \text{diag}(\Lambda_k)$. Now, minimizing the 2-norm of the difference of the two diagonal matrices in (28) is equivalent solving the problem: $\min\{\max_{ij} |(K - \sum_{k=1}^q \alpha_k \lambda_k \lambda_k^T)_{ij}| : \lambda_k \in \mathbb{R}^n, \alpha_k \in \mathbb{R}, k = 1, \dots, q\}$. Unfortunately, this problem cannot be solved easily. However, the variant that is based on minimizing the upper bound $\|K - \sum_{k=1}^q \alpha_k \lambda_k \lambda_k^T\|_2$ on its objective function can easily be solved. That is, we consider the following approximation problem for (28):

$$\min_{\alpha_k \in \mathbb{R}, \lambda_k \in \mathbb{R}^n, k=1, \dots, q} \left\| K - \sum_{k=1}^q \alpha_k \lambda_k \lambda_k^T \right\|_2, \quad (30)$$

The matrix 2-norm is chosen because this problem admits an analytical solution given by

$$\alpha_k = \sigma_k, \quad \lambda_k = u_k, \quad k = 1, \dots, q, \quad (31)$$

where σ_k is the k th largest eigenvalue (in absolute value) of the matrix K and u_k is the corresponding unit eigenvector vector.

Theorem 5.1 *For the preconditioner \widehat{M} given in (29) with the α_k 's and λ_k 's given by (31), we have*

$$\|M - \widehat{M}\|_2 \leq \|\widetilde{\mathcal{A}}\|_2^2 |\sigma_{q+1}|,$$

where σ_{q+1} is the $(q+1)$ -st largest eigenvalue of the matrix K . In case P and D were computed by part (a) of Lemma 5.1, $\|\tilde{\mathcal{A}}\|_2^2 = \|\mathcal{A}(W \otimes W)\mathcal{A}^T\|_2$, while if part (b) was used, $\|\tilde{\mathcal{A}}\|_2^2 = \|\mathcal{A}(U^{-1} \otimes U^{-1})\mathcal{A}^T\|_2$.

Proof. It is readily shown that

$$\begin{aligned} \|M - \widehat{M}\|_2 &= \left\| \tilde{\mathcal{A}} \left((\mathcal{I} + D \otimes D)^{-1} - \sum_{k=1}^q \sigma_k \text{Diag}(u_k) \otimes \text{Diag}(u_k) \right) \tilde{\mathcal{A}}^T \right\|_2 \\ &\leq \|\tilde{\mathcal{A}}\|_2^2 \left\| (\mathcal{I} + D \otimes D)^{-1} - \sum_{k=1}^q \sigma_k \text{Diag}(u_k) \otimes \text{Diag}(u_k) \right\|_2 \\ &\leq \|\tilde{\mathcal{A}}\|_2^2 \left\| K - \sum_{k=1}^q \sigma_k u_k u_k^T \right\|_2. \end{aligned}$$

Since $\sum_{k=1}^q \sigma_k u_k u_k^T$ is the best rank q approximation of K , the second norm in the last line above is given by σ_{q+1} ; see [13]. The last part follows since $\|\tilde{\mathcal{A}}\|_2^2 = \|\tilde{\mathcal{A}}\tilde{\mathcal{A}}^T\|_2 = \|\mathcal{A}(PP^T \otimes PP^T)\mathcal{A}^T\|_2$, using the form of PP^T given at the beginning of the subsection.

Remark 5.2 (a) While the Schur complement matrix (24) is positive definite, the matrix \widehat{M} may not be positive definite.

(b) For the numerical experiments in Section 7, we take q in Theorem 5.1 as follows:

$$q = \min \{ 15, \min \{ k : |\sigma_{k+1}| \leq 10^{-8} |\sigma_1| \} \}.$$

(c) Note that to solve (30), the easiest (though not necessarily the cheapest) mean is to compute the full eigenvalue decomposition of the $n \times n$ symmetric matrix K . If one is interested only in approximating K by the sum of a few rank-one matrices, then one can use variants of the Lanczos method to compute a partial eigenvalue decomposition of K .

(d) The construction of \widehat{M} can be made more efficient than computing each constituent term $\mathcal{A}(V_k \otimes V_k)\mathcal{A}^T$ separately. For example, the inner product with A_i in $\widehat{M}_{ij} = A_i \bullet (\sum_{k=1}^q \alpha_k V_k A_j V_k)$ need only be done once instead of q times. However, in the interest of keeping our implementation simple, we did not optimize the computational efficiency of \widehat{M} in Section 7.

5.2 Preconditioners for M , part II

For the special convex quadratic SDP with $Q = U \otimes U$, the middle operator in the Schur complement matrix M involves inverting an operator of the form $H = G_1 \otimes G_1 + G_2 \otimes G_2$, where G_1, G_2 are given symmetric positive definite matrices. Given that it is easy to invert an operator of the form $V \otimes V$, it is natural for us to consider approximating a

sum of symmetrized Kronecker products by a single term. Recall that a symmetrized Kronecker product $U \circledast U$ is an operator on \mathcal{S}^n , but it has a matrix representation using the operation \mathbf{svec} . Let $\text{Mat}(U \circledast U)$ in $\mathcal{S}^{\bar{n}}$ be defined by $\text{Mat}(U \circledast U)\mathbf{svec}(V) = \mathbf{svec}(U \circledast U(V))$. Then our problem is

$$\min_{V \in \mathcal{S}^n} \left\| \sum_{j=1}^2 \text{Mat}(G_j \circledast G_j) - \text{Mat}(V \circledast V) \right\|_F^2. \quad (32)$$

The above problem can be replaced by a simpler one which we will derive next. By noting that a symmetrized Kronecker product matrix $\text{Mat}(G \circledast G)$ is related to a standard Kronecker product by the formula $\text{Mat}(G \circledast G) = \Pi^T(G \otimes G)\Pi$, where the constant matrix $\Pi \in \mathbb{R}^{n^2 \times \bar{n}}$ has orthonormal columns (see the Appendix of [36]), we have

$$\begin{aligned} \left\| \sum_{j=1}^2 \text{Mat}(G_j \circledast G_j) - \text{Mat}(V \circledast V) \right\|_F &= \left\| \Pi^T \left(\sum_{j=1}^2 G_j \otimes G_j - V \otimes V \right) \Pi \right\|_F \\ &\leq \|\Pi\|_2^2 \left\| \sum_{j=1}^2 G_j \otimes G_j - V \otimes V \right\|_F. \end{aligned}$$

Note that $\|\Pi\|_2 = 1$. Thus instead of (32), we can consider solving the following problem:

$$\min_{V \in \mathcal{S}^n} \left\| \sum_{j=1}^2 G_j \otimes G_j - V \otimes V \right\|_F^2. \quad (33)$$

The problem (33) is a special case of a more general problem studied in [25] for approximating a sum of Kronecker products by a single Kronecker product, namely,

$$\min_{U, V \in \mathbb{R}^{n \times n}} \left\| \sum_{j=1}^q G_j \otimes K_j - U \otimes V \right\|_F^2, \quad (34)$$

where G_j, K_j are $n \times n$ matrices (not necessarily symmetric positive definite). It is shown in [25] that the optimal solution of (34) takes the form $U = \sum_{j=1}^q \alpha_j G_j$ and $V = \sum_{j=1}^q \beta_j K_j$, with the coefficients α_j, β_j being the optimal solution of the following equivalent non-convex minimization problem:

$$\min_{\alpha, \beta \in \mathbb{R}^q} f(\alpha, \beta) := \text{Tr}(\mathbf{G}\mathbf{K}) - 2\alpha^T \mathbf{G}\mathbf{K}\beta + (\alpha^T \mathbf{G}\alpha)(\beta^T \mathbf{K}\beta), \quad (35)$$

where $\mathbf{G}, \mathbf{K} \in \mathcal{S}_+^q$ are defined by $\mathbf{G}_{ij} = \text{Tr}(G_i^T G_j)$, $\mathbf{K}_{ij} = \text{Tr}(K_i^T K_j)$. For simplicity, we assume that $\{G_1, \dots, G_q\}$ and $\{K_1, \dots, K_q\}$ are linearly independent sets. Under this assumption, \mathbf{G} and \mathbf{K} are positive definite.

In [25], the optimal coefficients α and β in (35) are found by using an optimization software such as multilevel coordinate search. Here we show that the optimal solution can instead be found analytically.

Proposition 5.1 *The optimal objective value of (35) is given by*

$$f(\alpha^*, \beta^*) = \text{Tr}(\mathbf{G}\mathbf{K}) - \lambda_{\max}(\mathbf{G}\mathbf{K}),$$

where $\lambda_{\max}(\mathbf{G}\mathbf{K})$ is the largest eigenvalue of $\mathbf{G}\mathbf{K}$; (α^*, β^*) is a corresponding left and right eigenvector pair of $\mathbf{G}\mathbf{K}$. Let $R = \sum_{j=1}^q G_j \otimes K_j$. We have the following inequality for the relative error:

$$\frac{f(\alpha^*, \beta^*)}{\|R\|_F^2} = 1 - \frac{\lambda_{\max}(\mathbf{G}\mathbf{K})}{\text{Tr}(\mathbf{G}\mathbf{K})} \leq 1 - \frac{1}{q}.$$

Proof. We observe that f is coercive (so it has a minimizer) and that

$$\nabla f = 2 \begin{bmatrix} (\beta^T \mathbf{K} \beta) \mathbf{G} \alpha - \mathbf{G} \mathbf{K} \beta \\ (\alpha^T \mathbf{G} \alpha) \mathbf{K} \beta - \mathbf{K} \mathbf{G} \alpha \end{bmatrix}.$$

Thus the critical points of $f(\alpha, \beta)$ are given by the solutions of the following equations:

$$\mathbf{K} \beta = (\beta^T \mathbf{K} \beta) \alpha, \quad \mathbf{G} \alpha = (\alpha^T \mathbf{G} \alpha) \beta,$$

after making use of the fact that \mathbf{G} and \mathbf{K} are nonsingular. Now it is easy to see from the above equations that

$$\mathbf{K} \mathbf{G} \alpha = (\alpha^T \mathbf{G} \alpha) (\beta^T \mathbf{K} \beta) \alpha, \quad \mathbf{G} \mathbf{K} \beta = (\alpha^T \mathbf{G} \alpha) (\beta^T \mathbf{K} \beta) \beta. \quad (36)$$

This shows that the critical points $(\bar{\alpha}, \bar{\beta})$ are the left and right eigenvector pairs of $\mathbf{G}\mathbf{K}$. The corresponding objective value can be shown to be given by

$$f(\bar{\alpha}, \bar{\beta}) = \text{Tr}(\mathbf{G}\mathbf{K}) - (\bar{\alpha}^T \mathbf{G} \bar{\alpha}) (\bar{\beta}^T \mathbf{K} \bar{\beta}),$$

where the term $(\bar{\alpha}^T \mathbf{G} \bar{\alpha}) (\bar{\beta}^T \mathbf{K} \bar{\beta})$ is an eigenvalue of $\mathbf{G}\mathbf{K}$. Note that the eigenvalues of $\mathbf{G}\mathbf{K}$ are all real and non-negative. We have therefore shown that if $(\bar{\alpha}, \bar{\beta})$ is a critical point of f , then it is a left and right eigenvector pair of $\mathbf{G}\mathbf{K}$, and if $\lambda_i(\mathbf{G}\mathbf{K})$ is the corresponding eigenvalue, then $f(\bar{\alpha}, \bar{\beta}) = \text{Tr}(\mathbf{G}\mathbf{K}) - \lambda_i(\mathbf{G}\mathbf{K})$. It follows that $\min_{\alpha, \beta} f(\alpha, \beta) \geq \text{Tr}(\mathbf{G}\mathbf{K}) - \lambda_{\max}(\mathbf{G}\mathbf{K})$.

Now let $\hat{\lambda} = \lambda_{\max}(\mathbf{G}\mathbf{K})$ and $\hat{\beta}$ be the right eigenvector of $\mathbf{G}\mathbf{K}$ corresponding to $\hat{\lambda}$. Since $\mathbf{K} \succ 0$, we can define $\hat{\alpha} = \mathbf{K} \hat{\beta} / (\hat{\beta}^T \mathbf{K} \hat{\beta})$. Indeed, we shall now show that $(\hat{\alpha}, \hat{\beta})$ is a critical point corresponding to $\hat{\lambda}$. Note that $\hat{\beta}^T \hat{\alpha} = 1$. By the definition of $\hat{\alpha}$, it is clear that $\mathbf{K} \hat{\beta} = (\hat{\beta}^T \mathbf{K} \hat{\beta}) \hat{\alpha}$. Since $\mathbf{G} \hat{\alpha} = \mathbf{G} \mathbf{K} \hat{\beta} / (\hat{\beta}^T \mathbf{K} \hat{\beta}) = \hat{\lambda} \hat{\beta} / (\hat{\beta}^T \mathbf{K} \hat{\beta})$, we have that $\hat{\alpha}^T \mathbf{G} \hat{\alpha} = \hat{\lambda} \hat{\alpha}^T \hat{\beta} / (\hat{\beta}^T \mathbf{K} \hat{\beta}) = \hat{\lambda} / (\hat{\beta}^T \mathbf{K} \hat{\beta})$. Hence $\mathbf{G} \hat{\alpha} = (\hat{\alpha}^T \mathbf{G} \hat{\alpha}) \hat{\beta}$. Furthermore, $f(\hat{\alpha}, \hat{\beta}) = \text{Tr}(\mathbf{G}\mathbf{K}) - \hat{\lambda}$.

Now we prove the second statement. Note that it is easy to see that $\text{Tr}((G_j \otimes K_j)^T (G_l \otimes K_l)) = \text{Tr}(G_j^T G_l) \text{Tr}(K_j^T K_l)$. Thus we have $\|R\|_F^2 = \text{Tr}(\mathbf{G}\mathbf{K}) = \left\| \mathbf{G}^{1/2} \mathbf{K}^{1/2} \right\|_F^2$. We also have $\lambda_{\max}(\mathbf{G}\mathbf{K}) = \left\| \mathbf{G}^{1/2} \mathbf{K}^{1/2} \right\|_2^2$. By a standard inequality between the matrix 2-norm and Frobenius norm, the required inequality is established. \square

Although it seems plausible that the optimal solution of (34) could very well also be the optimal solution of (32), we note that simple examples demonstrate that this is not necessarily true.

Proposition 5.2 *Suppose G_j and K_j are symmetric positive definite for $j = 1, \dots, q$. Then the optimal solution pair U^* and V^* in (34) are also symmetric positive definite.*

Proof. Since $G_j, K_j, j = 1, \dots, q$ are positive definite, \mathbf{G} and \mathbf{K} are non-negative matrices. This implies that $\mathbf{G}\mathbf{K}$ is a non-negative matrix. By the Perron-Frobenius Theorem [19, p. 500], $\lambda_{\max}(\mathbf{G}\mathbf{K})$ is algebraically simple and there exists a corresponding left and right eigenvector pair α^*, β^* for which the vectors are non-negative. In this case, the matrices $U^* = \sum_{j=1}^q \alpha_j^* G_j, V^* = \sum_{j=1}^q \beta_j K_j$ are symmetric positive definite. \square

For the problem (33), the optimal solution is given by $V^* = \sum_{j=1}^2 \alpha_j^* G_j$, with α^* being a unit eigenvector corresponding to the largest eigenvalue of \mathbf{G}^2 . Note that by Proposition 5.2, V^* is symmetric positive definite. Since $V^* \otimes V^*$ is an approximation of \mathcal{H} , we can naturally precondition the Schur complement matrix in (24) using the following matrix:

$$\mathcal{A}[(V^*)^{-1} \otimes (V^*)^{-1}] \mathcal{A}^T. \quad (37)$$

While the preconditioner we have constructed in (29) may not be positive definite, the above preconditioner is guaranteed to be positive definite.

In the numerical experiments in Section 7, we found that the preconditioner constructed by approximating the term $I \otimes I + D \otimes D$ in (24) by a single symmetrized Kronecker product of the form $\Sigma \otimes \Sigma$ (where Σ is a linear combination of I and D obtained by solving (33) with $G_1 = I$ and $G_2 = D$) is more effective than the one constructed from approximating $\mathcal{H} = W^{-1} \otimes W^{-1} + U \otimes U$ directly. In this case, the preconditioner is given by

$$\widehat{M} = \mathcal{A}[(P\Sigma^{-1}P^T) \otimes (P\Sigma^{-1}P^T)] \mathcal{A}^T. \quad (38)$$

We will use the above preconditioner in Section 7.

6 Linear semidefinite programming with a simple upper bound

Consider the following semidefinite program with a simple upper bound:

$$\begin{aligned} \min_X \quad & C \bullet X \\ \mathcal{A}(X) = b, \quad & 0 \preceq X \preceq U, \end{aligned} \quad (39)$$

where $U \succeq 0$ is a given matrix. An example of (39) comes from minimizing the sum of the largest q eigenvalues of an affine function of symmetric matrices, namely,

$$\min \left\{ \sum_{k=1}^q \lambda_k(\mathcal{A}^T y - C) : y \in \mathbb{R}^m \right\}. \quad (40)$$

In [1, (4.7)], it is shown that (40) is equivalent to a linear SDP with a simple upper bound:

$$\min \{C \bullet X : \mathcal{A}(X) = 0, I \bullet X = q, 0 \preceq X \preceq I\}. \quad (41)$$

To derive the KKT conditions for (39), it is convenient to express (39) in the following standard form:

$$\min_X \left\{ C \bullet X : \begin{bmatrix} \mathcal{A} \\ -I \end{bmatrix} X + \begin{bmatrix} 0 \\ -I \end{bmatrix} V = \begin{bmatrix} b \\ -U \end{bmatrix}, X, V \succeq 0 \right\}. \quad (42)$$

We see that converting the problem (39) to the standard form introduces \bar{n} extra equality constraints in the primal problem. Thus it is extremely expensive to solve (39) by treating it as a standard linear SDP.

The dual problem corresponding to (42) is given by

$$\max \{b^T y - U \bullet Z : \mathcal{A}^T y - Z + S = C, S, Z \succeq 0\}. \quad (43)$$

The perturbed KKT conditions for (42) and (43) are given by

$$\begin{aligned} \mathcal{A}X &= b \\ X + V &= U \\ \mathcal{A}^T y - Z + S &= C \\ XS &= \nu I \\ VZ &= \nu I. \end{aligned} \quad (44)$$

The symmetrized Newton equation corresponding to the above system is given by

$$\begin{aligned} \mathcal{A}\Delta X &= r^p := b - \mathcal{A}X \\ \Delta X + \Delta V &= R^u := U - X - V \\ \mathcal{A}^T \Delta y - \Delta Z + \Delta S &= R^d := C - \mathcal{A}^T y + Z - S \\ \mathcal{E}_1 \Delta X + \mathcal{F}_1 \Delta S &= R_1^c := \sigma \mu I - H_1(XS) \\ \mathcal{E}_2 \Delta V + \mathcal{F}_2 \Delta Z &= R_2^c := \sigma \mu I - H_2(VZ), \end{aligned} \quad (45)$$

where $\mathcal{E}_1, \mathcal{F}_1$ and $\mathcal{E}_2, \mathcal{F}_2$ are linear operators in \mathcal{S}^n that depend on the symmetrization schemes chosen for X, S and V, Z , respectively. It is readily shown that the search direction $(\Delta X, \Delta V, \Delta y, \Delta S, \Delta Z)$ corresponding to the Newton equation (45) can be computed by solving the following linear system:

$$\begin{bmatrix} -(\mathcal{F}_1^{-1} \mathcal{E}_1 + \mathcal{F}_2^{-1} \mathcal{E}_2) & \mathcal{A}^T \\ \mathcal{A} & 0 \end{bmatrix} \begin{bmatrix} \Delta X \\ \Delta y \end{bmatrix} = \begin{bmatrix} h \\ r^p \end{bmatrix}$$

where

$$h = R^d - \mathcal{F}_1^{-1} R_1^c + \mathcal{F}_2^{-1} R_2^c - \mathcal{F}_2^{-1} \mathcal{E}_2 R^u.$$

By eliminating ΔX from the augmented equation above, we get the following Schur complement equation:

$$\mathcal{A}(\mathcal{F}_1^{-1}\mathcal{E}_1 + \mathcal{F}_2^{-1}\mathcal{E}_2)^{-1} \mathcal{A}^T \Delta y = r^p + \mathcal{A}(\mathcal{F}_1^{-1}\mathcal{E}_1 + \mathcal{F}_2^{-1}\mathcal{E}_2)^{-1} h. \quad (46)$$

The operator $\mathcal{F}_1^{-1}\mathcal{E}_1 + \mathcal{F}_2^{-1}\mathcal{E}_2$ in (46) is costly to invert in general. Hence constructing the Schur complement matrix in (46) can be very expensive. The reader may recall that in contrast, for the case of linear programming, having an upper bound vector does not add any extra computational cost because the middle matrix in (46) is the sum of two diagonal matrices.

The middle matrix in (46) has exactly the same structure as that in (22) if the NT scalings are chosen. In this case, $\mathcal{F}_1^{-1}\mathcal{E}_1 + \mathcal{F}_2^{-1}\mathcal{E}_2 = W_1^{-1} \otimes W_1^{-1} + W_2^{-1} \otimes W_2^{-1}$. Thus the Schur complement matrix in (46) can be reduced to exactly the same form as the matrix M in (24), and the preconditioners proposed in the last section can be used when solving (46) by an iterative solver.

The problem (39) is an example of a problem where the decision variables are constrained to lie in the intersection of two cones. For such problems, the Schur complement equation arising from an interior-point iteration will have the form (46). Andersen studies such structures for the intersection of a linear and a second-order cone [5]. Another example appears in [15] where the authors study a problem very similar to (39). This problem comes from a robust optimization formulation of a convex quadratic programming problem and the upper bound is expressed using componentwise inequalities rather than the semidefinite inequality. In this case, the matrix $\mathcal{F}_2^{-1}\mathcal{E}_2$ is diagonal and the situation is similar to the QSDP problem with a diagonal Q operator.

7 Numerical experiments

To evaluate the performance of our interior-point algorithm, we consider the following classes of test problems:

- E1. Quadratic SDPs arising from the nearest correlation matrix problem where $Q(X) = X$. We use the linear map $\mathcal{A}(X) = \text{diag}(X)$, and b is the vector of all ones; see (26). We generate the matrix $-C$ in the same way as Higham did in [18, p.340]. It is generated from the MATLAB function `gallery('randcorr', ...)` with a random symmetric perturbation of Frobenius norm 10^{-4} added.
- E2. Same as E1 but the matrix C is generated as follows: `T = 2*rand(n)-1; C = -0.5*(T+T')`; such a matrix is considered in the numerical experiments in [28].
- E3. Same as E1 but the matrix C is generated as in [28] as follows: `T = [ones(n/2), zeros(n/2); zeros(n/2), eye(n/2)]; C = -T - 1e4*diag(2*rand(n,1)-1)`;
- E4. Same as E1 but $Q(X) = UXU$, with $U \succ 0$ generated randomly as follows: `[Q,R] = qr(randn(n)); beta = 10^(-4/(n-1)); U = Q*diag(beta.^[0:n-1])*Q'`. The matrix C in E1 is replaced by UCU . Note that in this case, the condition number of Q is 10^8 .

- E5. Same as E2 but $\mathcal{Q}(X) = UXU$, with U generated as in E4. The matrix C in E2 is replaced by UCU .
- E6. Same as E3 but $\mathcal{Q}(X) = UXU$, with U generated as in E4. The matrix C in E3 is replaced by UCU .
- E7. Same as E1 but $\mathcal{Q}(X) = UXU$, with U generated as follows:
 $\text{beta} = 10^{(2/(n-1))}$; $U = \text{diag}(\text{beta}.^{[0:n-1]})$. The matrix C in E1 is replaced by UCU . In this case, the condition number of \mathcal{Q} is 10^4 .
- E8. Same as E2 but $\mathcal{Q}(X) = UXU$, with U generated as in E7. The matrix C in E2 is replaced by UCU .
- E9. Same as E3 but $\mathcal{Q}(X) = UXU$, with U generated as in E7. The matrix C in E3 is replaced by UCU .
- E10. Linear SDPs with simple upper bounds (41) arising from (40). Here we take $q = 5$, the matrix C is generated randomly in MATLAB as follows: $T = \text{randn}(n)$; $T = 0.5*(T+T')$; $C = T + \text{norm}(T,2)*I$; the linear map \mathcal{A} is chosen to be $\mathcal{A}(X) = [A_1 \bullet X, \dots, A_{n-1} \bullet X]$, where $A_k = e_k e_{k+1}^T + e_{k+1} e_k^T$ for $k = 1, \dots, n-1$.

We use 4 variants of Algorithm IP-QSDP to solve each test problem, namely,

- A1. Algorithm IP-QSDP with search direction computed via (13) and (24) by a direct solver;
- A2. Algorithm IP-QSDP with search direction computed via (13) and (24) by PSQMR with no preconditioning.
- A3. Algorithm IP-QSDP with search direction computed via (13) and (24) by PSQMR with a hybrid preconditioner chosen as follows: it is taken to be the preconditioner (29) if it is positive definite; otherwise, it is taken to be the preconditioner (38).
- A4. Algorithm IP-QSDP with search direction computed via (13) and (24) by PSQMR with preconditioner (38).

We implemented the algorithms in MATLAB (version 7.0) and the experiments were conducted on a Pentium IV 3.0GHz PC with 2GB of RAM. We stopped the algorithms when the accuracy measure ϕ in (8) was less than 10^{-7} , or when the algorithms did not improve both the duality gap and infeasibilities. The stopping criterion used to solve the system (13) is described in the paragraph just before Section 4.1. We also set the maximum number of PSQMR steps allowed per solve to n .

The initial iterate for all the algorithms was taken to be the default starting point in [37]. For the problem sets E1–E6, it is given by

$$X^0 = \frac{n}{\sqrt{2}}I, \quad y^0 = 0, \quad S^0 = \sqrt{n}I.$$

The performance results of our algorithms are given in Table 1. The columns corresponding to “it” give the total number of interior-point iterations required to solve each problem, whereas the columns “psqmr” give the average number of PSQMR steps required to solve each of the two linear systems (24) during the computation of the predictor and corrector directions at each interior-point iteration. Note that we did

not run Algorithm A1 for some of the larger problems and the corresponding entries in Table 1 are left blank.

It is worth making some observations we may derive from the performance table.

1. Solving (13) via a direct solver is extremely expensive. For the problem in E1-1600, it is at least 20 times more expensive than the algorithms using iterative solvers to solve (13).
2. Based on the stopping criterion we proposed in Section 4, the algorithms that use an iterative method to solve (13) took about the same number of interior-point iterations to converge compared to the algorithm using a direct method. This indicates that the inexact search directions are computed to sufficient accuracy, and thus the residual errors do not degrade the outer iterations.
3. The test examples considered in [18] and [28] for the unweighted nearest correlation matrix problem are easy problems that can be solved by an iterative solver even without preconditioning. As we can observe from the test problems in E1–E3, the SQMR method takes an average of 2 to 4 steps to solve the Schur complement equation (13) without any preconditioning. This indicates that the coefficient matrix M is very well-conditioned throughout the entire course of interior-point iterations. With such a well-conditioned system, the preconditioners proposed in (29) and (38) cannot offer any saving in the computation time because of the overheads involved in their construction. For these easy problems, the condition number of M stays bounded even when the duality gap decreases to zero. This possibility is consistent with Corollary 4.1.
4. The conditioning of the matrix M becomes slightly worse for the weighted nearest correlation matrix problems considered in E4 and E5. This can be seen from the slight increase in the average number of SQMR steps required to solve (13). The preconditioned systems generally take fewer steps to converge, but the reduction in the number of PSQMR steps is not enough to offset the preconditioning overhead.
5. The test problems in E6-E10 truly demonstrate the effectiveness of the preconditioners (29) and (38). For the test problem in E7-2000, the SQMR method takes an average of 106.0 steps to solve (13) whereas the corresponding numbers for the system preconditioned by (29) and (38) are 1.0 and 3.1, respectively.
6. The number of interior-point iterations required by our proposed primal-dual interior-point method grows very modestly with the problem dimension n . In all the test problems, the number of iterations required is less than 30. In contrast, for problems similar to those in E3 considered in [28], the number of iterations required to solve the Lagrangian dual problems is about 350.
7. If we compare the CPU time taken by Algorithm A2 to solve the problem E1-2000 with that in [28], each interior-point iteration of our algorithm is about seven times slower (after taking into account the difference in the speed of the machines used). Such a big discrepancy can be accounted for by the overhead incurred at each IPM iteration. First note that each iteration in [28] is a quasi-Newton (BFGS) iteration for which the dominant cost is in computing a full

eigenvalue decomposition of the primal variable at a cost of $26n^3/3$ flops. Notice that for this problem, only one PSQMR step on the average is required to solve the system (13) each for the predictor and corrector steps at each IPM iteration, and this accounts for a total of only $8n^3$ flops. Besides the cost of solving (13), each of our IPM iteration incurs a large overhead in computing the right-hand side vector h in (13) and ΔX and ΔS from (16). The total overhead accrued from such computations for both the predictor and corrector steps is about $28n^3$ flops. Furthermore, each of our IPM iterations also requires two full eigenvalue decompositions in computing the NT scaling matrix W and the decomposition of \mathcal{H}^{-1} in Lemma 5.1, and this contributes another $56n^3/3$ flops to the overhead. Thus we can see that the total overhead cost at each IPM iteration is at least about $140n^3/3$ flops. This alone is already more than five times the cost of each BFGS iteration in [28]. However, our algorithm seems to require only about half as many iterations as Malick's for problems of the class E2, and only about a twentieth for problems of class E3.

8. It is clear from (29) and (38) that the overhead incurred in constructing the first preconditioner will be more than that for the second. This is reflected in the CPU times in Table 1. For example, Algorithm A4 takes less time to solve E5-2000 (similarly for E10-2000) than Algorithm A3, even though the former required more PSQMR steps per solve. Generally, both preconditioners are quite effective on all the problem classes considered. However, the first preconditioner is much more effective than the second for the problem class E6 and for the smaller problems in class E10 (we do not know why the larger problems in this class seem much easier for the second preconditioner to handle).

8 Conclusion and future research

We considered a primal-dual path-following Mehrotra-type predictor-corrector method for solving convex quadratic SDP problems. For the special case when the derivative of the quadratic term $\mathcal{Q}(X)$ has the form UXU , we suggested computing the search direction at each iteration based on the Schur complement equation, using the PSQMR iterative solver with one of two appropriately constructed preconditioners. Numerical experiments on a variety of QSDPs with matrices of dimension up to 2000 showed that our methods are quite efficient and robust. We also extended our methods to solving linear SDP problems with upper bound constraints on the primal matrix variable.

Thus, using our preconditioners, certain classes of quadratic SDPs can be solved at a cost not much greater than linear SDPs of the same size using primal-dual interior-point methods. However, such methods have an inherently higher complexity than first-order methods such as the spectral bundle method of Helmsberg and Rendl [16] and the nonlinear-programming based method of Burer, Monteiro, and Zhang [9], so it is of interest to ask whether such methods can be extended to quadratic SDPs. We note that in fact the spectral bundle method solves a (very small) quadratic SDP using a primal-dual interior-point method at every iteration to determine its search direction.

We observed from our computations that the matrix \widehat{M} in (29) is usually an ex-

tremely good approximation to M for the value q selected based on Remark 5.2(b). An interesting idea that is worth exploring is to use the direction Δy computed from $\widehat{M}\Delta y = h$ for each IPM iteration, instead of the direction computed from (13) itself. In that case, Δy can be computed from the approximate system straightforwardly by a direct solver.

As mentioned in the introduction, our ultimate goal is to solve problems with a general positive semidefinite Q , in which case the search direction at each IPM iteration has to be computed from the augmented equation (11). Our hope is to solve this equation efficiently via an iterative solver with appropriately constructed preconditioners.

Through the years and in the current paper, our interest in the robust and efficient implementation of interior-point methods for various classes of conic optimization problems and our continued work on SDPT3 were motivated, and in many cases inspired, by the friendly competition we received from the SDPA software developed by Kojima's group. For example, the technique on exploiting sparsity in SDPT3 is largely based on the work by Fujisawa, Kojima and Nakata [12]. We dedicate this work to Masakazu Kojima on the occasion of his 60th birthday.

Acknowledgements

We would like to thank an anonymous referee for his very helpful comments on an earlier version of this paper.

Table 1: Performance of the Algorithms A1–A4 on the problem sets E1–E10.

	A1				A2				A3				A4			
n	it	ϕ	Time (s)		it	ϕ	Time (s)	psqmr	it	ϕ	Time (s)	psqmr	it	ϕ	Time (s)	psqmr
E1	200	13	3.7e-08	34.1	13	3.6e-08	8.0	3.5	13	3.7e-08	10.4	1.0	13	3.7e-08	9.0	3.3
	400	13	6.3e-08	281.2	13	6.3e-08	38.3	2.9	13	6.3e-08	47.3	1.0	13	6.1e-08	39.4	2.8
	800	14	4.0e-08	3348.9	14	4.0e-08	269.4	2.8	14	4.0e-08	322.0	1.0	14	4.0e-08	277.7	2.7
	1600	14	3.3e-08	43273.4	14	3.3e-08	1730.4	2.5	14	3.3e-08	2245.7	1.0	14	3.3e-08	1747.7	2.5
	2000				14	5.1e-08	3920.1	2.4	14	5.1e-08	4266.2	1.0	14	5.1e-08	4075.9	2.4
E2	200	9	2.2e-08	22.9	9	2.2e-08	5.1	2.4	9	2.2e-08	6.5	1.6	9	2.2e-08	5.1	2.0
	400	10	2.5e-08	212.0	10	2.4e-08	28.1	2.4	10	2.5e-08	36.5	1.6	10	2.5e-08	28.4	2.0
	800	10	2.0e-08	2321.7	10	2.1e-08	172.9	2.1	10	2.0e-08	236.8	1.5	10	1.9e-08	175.8	1.8
	1600				10	9.0e-08	1143.3	1.9	10	3.6e-08	1646.7	1.4	10	3.6e-08	1162.4	1.6
	2000				10	7.8e-08	2735.4	1.8	10	3.0e-08	3093.6	1.2	10	3.0e-08	2812.2	1.6
E3	200	12	1.8e-08	26.5	12	2.5e-08	6.5	1.7	12	1.8e-08	7.7	1.0	12	4.6e-08	6.9	1.6
	400	12	1.6e-08	269.6	12	2.1e-08	37.6	1.6	12	1.6e-08	44.3	1.0	12	1.7e-08	39.1	1.6
	800	13	1.4e-08	3245.9	13	1.5e-08	269.0	1.7	13	5.2e-08	338.5	1.0	13	1.4e-08	276.7	1.7
	1600				14	6.8e-09	1932.0	1.7	14	6.7e-09	2302.3	1.1	14	6.7e-09	1974.7	1.7
	2000				14	1.1e-08	3582.9	1.6	14	1.0e-08	4347.9	1.1	14	1.0e-08	3761.2	1.7
E4	200	10	3.8e-08	26.4	10	3.1e-08	7.7	10.2	10	3.8e-08	6.9	1.0	10	5.2e-08	6.6	5.8
	400	11	8.3e-08	257.6	11	6.2e-08	44.4	7.9	11	8.3e-08	46.1	1.5	11	9.9e-08	41.5	5.8
	800	14	1.5e-08	3663.5	14	6.4e-08	349.4	6.6	14	1.5e-08	396.2	2.8	15	9.1e-09	372.4	5.9
	1600				14	1.9e-08	2191.8	4.7	13	7.6e-08	2514.3	1.2	14	1.2e-08	2255.7	4.6
	2000				14	7.4e-08	4067.1	4.4	14	1.7e-08	5451.1	3.0	14	4.3e-08	4160.6	4.3
E5	200	12	2.6e-08	32.6	12	3.1e-08	10.2	11.7	12	2.6e-08	8.9	1.2	12	2.7e-08	7.7	4.7
	400	12	4.2e-08	281.3	12	4.4e-08	48.5	7.7	12	4.2e-08	50.9	1.2	12	4.1e-08	43.7	4.8
	800	13	2.2e-08	3358.3	13	2.2e-08	319.3	6.4	13	2.2e-08	361.9	2.5	13	2.1e-08	296.2	4.3
	1600				14	2.6e-08	2237.2	5.1	14	9.8e-09	2702.5	2.8	14	9.1e-09	2383.5	4.2
	2000				13	6.9e-08	3769.9	4.4	13	3.0e-08	4940.3	2.4	13	3.6e-08	3792.4	3.9

Table 1: Performance of the Algorithms A1–A4 on the problem sets E1–E10.

	A1				A2				A3				A4			
n	it	ϕ	Time (s)		it	ϕ	Time (s)	psqmr	it	ϕ	Time (s)	psqmr	it	ϕ	Time (s)	psqmr
E6	200	20	1.2e-08	54.2	21	9.6e-09	32.1	32.9	20	1.2e-08	13.4	1.1	20	1.3e-08	19.0	14.7
	400	21	1.4e-08	487.7	23	2.5e-08	218.3	39.5	21	1.4e-08	76.7	1.0	21	4.2e-08	111.0	15.4
	800	22	5.5e-08	5633.7	24	4.9e-08	1544.3	43.7	22	5.5e-08	522.6	1.0	22	7.3e-08	696.0	13.6
	1600				27	8.6e-08	16579.6	63.9	25	2.4e-08	4223.8	1.0	25	9.6e-08	6759.0	20.6
	2000				28	8.0e-08	30566.7	62.1	26	1.1e-08	8632.0	1.0	27	7.1e-08	14097.6	21.7
E7	200	7	2.6e-09	18.5	8	5.5e-09	18.6	90.0	7	2.6e-09	4.5	1.0	8	2.7e-09	5.2	3.2
	400	7	1.3e-08	153.0	8	1.3e-08	121.5	97.1	7	1.3e-08	24.5	1.0	8	4.6e-09	23.5	3.0
	800	7	3.3e-08	1623.8	8	3.0e-08	860.9	102.1	7	3.3e-08	158.1	1.0	8	8.7e-09	151.7	3.1
	1600	7	3.7e-08	21708.1	8	3.4e-08	6919.5	105.9	7	3.7e-08	1159.0	1.0	8	2.1e-08	1072.6	3.1
	2000	7	5.2e-08	46041.5	8	4.8e-08	12652.1	106.0	7	5.2e-08	2140.6	1.0	8	1.2e-08	1993.8	3.1
E8	200	7	2.5e-08	18.6	9	1.3e-08	16.4	68.2	7	4.8e-08	5.2	1.1	8	6.7e-08	4.8	4.5
	400	7	6.7e-08	153.0	9	2.2e-08	100.9	65.4	8	7.9e-08	31.7	1.9	9	1.8e-08	29.3	4.8
	800	8	2.8e-08	1889.4	9	9.5e-08	604.9	57.2	9	5.4e-08	228.9	2.4	9	7.4e-08	187.1	5.2
	1600	8	8.9e-08	25585.9	10	3.7e-08	5021.4	50.2	10	3.7e-08	2038.2	3.1	10	3.8e-08	1662.5	5.8
	2000				10	4.8e-08	8590.9	48.7	10	4.4e-08	3441.8	3.1	10	4.6e-08	2850.4	5.7
E9	200	12	2.0e-08	34.5	13	4.8e-08	17.1	41.4	12	2.0e-08	8.3	1.0	13	1.3e-08	8.2	6.2
	400	13	1.8e-08	330.4	14	2.6e-08	119.8	45.9	13	1.8e-08	53.1	1.0	14	9.8e-09	53.9	7.2
	800	13	9.5e-08	3623.9	14	9.4e-08	815.6	45.4	13	9.5e-08	359.4	1.0	15	1.3e-08	401.7	8.2
	1600	15	2.5e-08	50155.2	16	2.1e-08	7243.8	46.9	15	2.5e-08	2964.0	1.0	16	2.0e-08	3185.9	9.7
	2000				16	2.7e-08	13428.5	47.3	15	4.6e-08	5549.2	1.0	17	1.0e-08	6379.2	10.2
E10	200	14	5.5e-08	55.2	15	6.1e-08	71.1	76.1	14	5.2e-08	18.5	1.0	14	8.2e-08	61.1	69.2
	400	14	7.1e-08	325.2	14	6.9e-08	335.8	96.0	14	7.5e-08	106.7	1.4	14	1.0e-07	313.7	66.6
	800	15	6.6e-08	3955.8	15	1.1e-07	4312.7	236.4	15	1.4e-07	508.8	1.3	15	9.1e-08	650.8	12.9
	1600	17	2.7e-07	56457.0	17	3.3e-07	62418.6	440.9	17	1.3e-07	5818.5	1.3	17	7.7e-08	3917.3	6.4
	2000				18	1.4e-06	136231.9	520.2	17	1.2e-07	9128.1	1.1	17	8.6e-08	7046.7	5.4

References

- [1] F. Alizadeh, *Interior point methods in semidefinite programming with applications to combinatorial optimization*, SIAM J. Optimization, 5 (1995), pp. 13–51.
- [2] F. Alizadeh, J.-P. A. Haeberly, and M. L. Overton, *Complementarity and nondegeneracy in semidefinite programming*, Math. Programming, 77 (1997), pp. 111–128.
- [3] F. Alizadeh, J.-P. A. Haeberly, and M. L. Overton, *Primal-dual interior-point methods for semidefinite programming: convergence results, stability and numerical results*, SIAM J. Optimization, 8 (1998), pp. 746–768.
- [4] A. Y. Alfakih, A. Khandani, and H. Wolkowicz, *Solving Euclidean distance matrix completion problems via semidefinite programming*, Computational Optimization and Applications, 12 (1999), pp. 13–30.
- [5] E. D. Andersen, *On exploiting structure induced when modelling an intersection of cones in conic optimization*, Research Report, July 2004. available from <http://www.optimization-online.org>.
- [6] M. F. Anjos, N. J. Higham, P. L. Takouda, and H. Wolkowicz, *A semidefinite programming approach for the nearest correlation matrix problem*, Research Report, Department of Combinatorics and Optimization, University of Waterloo, 2003.
- [7] P. Apkarian, D. Noll, J.-P. Thevenet, and H. D. Tuan, *A spectral quadratic-SDP method with applications to fixed-order H_2 and H_∞ synthesis*, Research Report, Universite Paul Sabatier, Toulouse, France, 2004.
- [8] S. Boyd and L. Xiao, *Least-squares covariance matrix adjustment*, preprint, Department of Electrical Engineering, Stanford University.
- [9] S. Burer, R. D. C. Monteiro, and Y. Zhang, *Solving a class of semidefinite programs via nonlinear programming*, Mathematical Programming, 93 (2002), pp. 97–122.
- [10] J. W. Demmel, *Applied Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [11] R. W. Freund and N. M. Nachtigal, *A new Krylov-subspace method for symmetric indefinite linear systems*, Proceedings of the 14th IMACS World Congress on Computational and Applied Mathematics, Atlanta, USA, W.F. Ames ed., July 1994, pp. 1253–1256.
- [12] K. Fujisawa, M. Kojima, and K. Nakata, *Exploiting sparsity in primal-dual interior-point method for semidefinite programming*, Mathematical Programming, 79 (1997), pp. 235–253.
- [13] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 2nd ed., Johns Hopkins University Press, Baltimore, MD, 1989.
- [14] M. Halicka, E. de Klerk, and C. Roos, *Limiting behaviour of the central path in semidefinite optimization*, Optimization Methods and Software, 20 (2005), pp. 99–113.

- [15] B. V. Halldórsson and R. H. Tütüncü, *An interior-point method for a class of saddle point problems*, Journal of Optimization Theory and Applications, 116 (2003) pp. 559–590.
- [16] C. Helmberg and F. Rendl, *A spectral bundle method for semidefinite programming*, SIAM Journal on Optimization, 10 (2000), pp. 673–696.
- [17] C. Helmberg, F. Rendl, R. Vanderbei, and H. Wolkowicz, *An interior-point method for semidefinite programming*, SIAM Journal on Optimization, 6 (1996), pp. 342–361.
- [18] N. J. Higham, *Computing the nearest correlation matrix — a problem from finance*, IMA J. Numerical Analysis, 22 (2002), pp. 329–343.
- [19] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, Cambridge, 1985.
- [20] C. Kanzow, C. Nagel, H. Kato, and M. Fukushima, *Successive linearization methods for nonlinear semidefinite programming*, preprint.
- [21] M. Kojima, S. Shindoh, and S. Hara, *Interior-point methods for the monotone linear complementarity problem in symmetric matrices*, SIAM J. Optimization, 7 (1997), pp. 86–125.
- [22] M. Kojima, M. Shida, and S. Shindoh, *Reduction of Monotone Linear Complementarity Problems over Cones to Linear Programs over Cones*, Acta Mathematica Vietnamica, 22 (1997), pp. 147–157.
- [23] M. Kojima, M. Shida, and S. Shindoh, *Local Convergence of Predictor-Corrector Infeasible-Interior-Point Method for SDPs and SDLCPs*, Math. Prog. 80 (1998) pp. 129–160.
- [24] M. Kojima, M. Shida, and S. Shindoh, *Search Directions in the SDP and the Monotone SDLCP: Generalization and Inexact Computation*, Math. Prog. 85 (1999) pp. 51–80.
- [25] A. N. Langville and W. J. Stewart, *A Kronecker product approximate preconditioner for SANs*, Numerical Linear Algebra with Applications, 11 (2004), pp. 723–752.
- [26] Z. Lu, R. D. C. Monteiro, and J. W. O’Neal, *An iterative solver-based infeasible primal-dual path-following algorithm for convex QP*, preprint, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta.
- [27] Z.-Q. Luo, J. F. Sturm, and S. Zhang, *Superlinear convergence of a symmetric primal-dual path following algorithm for semidefinite programming*, SIAM J. on Optimization, 8 (1998), pp. 59–81.
- [28] J. Malick, *A dual approach to semidefinite least-squares problems*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 272–284.
- [29] R. D. C. Monteiro and P. R. Zanjácomo, *Implementation of primal-dual methods for semidefinite programming based on Monteiro and Tsuchiya Newton directions and their variants*, Optimization Methods and Software, 11/12 (1999), pp. 91–140.

- [30] R. D. C. Monteiro, *Primal-dual path-following algorithms for semidefinite programming*, SIAM J. Optimization, 7 (1997), pp. 663–678.
- [31] R. D. C. Monteiro and Y. Zhang, *A unified analysis for a class of path-following primal-dual interior-point algorithms for semidefinite programming*, Math. Prog., 81 (1998), pp. 281–299.
- [32] J. W. Nie and Y. X. Yuan, *A predictor-corrector algorithm for QSDP combining Dikin-type and Newton centering steps*, Annals of Operations Research, 103 (2001), pp. 115–133.
- [33] H. Qi and D. Sun, *Quadratic convergence and numerical experiments of Newton’s method for computing the nearest correlation matrix*, preprint, Department of Mathematics, National University of Singapore.
- [34] Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Company, Boston, 1996.
- [35] J. F. Sturm, *Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones*, Optimization Methods and Software, 11 & 12 (1999), pp. 625–653.
- [36] M. J. Todd, K. C. Toh, and R. H. Tütüncü, *On the Nesterov-Todd direction in semidefinite programming*, SIAM J. Optimization, 8 (1998), pp. 769–796.
- [37] R. H. Tütüncü, K. C. Toh, and M. J. Todd, *Solving semidefinite-quadratic-linear programs using SDPT3*, Math. Prog. 95 (2003), pp. 189–217.
- [38] K. C. Toh, K. K. Phoon, and S. H. Chan, *Block preconditioners for symmetric indefinite linear systems*, Int. J. Numerical Methods in Engineering, 60 (2004), pp. 1361–1381.
- [39] K. C. Toh, *Some new search directions for primal-dual interior point methods in semidefinite programming*, SIAM J. Optimization, 11 (2000), pp. 223–242.