Solving second order cone programming via a reduced augmented system approach

Zhi Cai *and Kim-Chuan Toh †

Abstract

The standard Schur complement equation based implementation of interior-point methods for second order cone programming may encounter stability problems in the computation of search directions, and as a consequence, accurate approximate optimal solutions are sometimes not attainable. Based on the eigenvalue decomposition of the $(1,1)$ block of the augmented equation, a reduced augmented equation approach is proposed to ameliorate the stability problems. Numerical experiments show that the new approach can achieve much more accurate approximate optimal solutions than the Schur complement equation based approach.

Keywords: second order cone programming, augmented equation, Nesterov-Todd direction, stability

1 Introduction

A second order cone programming (SOCP) problem is a linear optimization problem over a cross product of second order convex cones. In [15], an extended list of application problems are shown to be SOCP problems. In particular, linear programming (LP) problem, and convex quadratically constrained quadratic programming (QCQP) are both subclasses of SOCP. SOCP itself is a subclass of semidefinite programming (SDP). SOCP has since found many more applications, notably in the area of filter design [5, 21] and in limit analysis of collapses of solid bodies [6]. For a comprehensive introduction to SOCP, we refer the reader to the paper by Alizadeh and Goldfarb [1].

*High Performance Computing for Engineered Systems (HPCES), Singapore-MIT Alliance, 4 Engineering Drive 3, Singapore 117576. (smap0035@nus.edu.sg).

†Department of Mathematics, National University of Singapore, 2 Science Drive 2, Singapore 117543, Singapore (mattohkc@math.nus.edu.sg).

In theory, SOCP problems can be solved as SDP problems. However it is far more efficient computationally to solve SOCP problems directly. Global polynomial convergence results of interior-point methods (IPMs) for SOCP have been well established; see [16] and the references therein. But there are few research works published on the implementation of IPMs for solving SOCP, or on methods for solving an SOCP efficiently and stably to obtain highly accurate approximate optimal solutions. The main objective of this paper is to propose a method that can solve an SOCP to high accuracy, but with comparable, or moderately higher cost than the standard IPMs employing the Schur complement equation (SCE) approach.

Given a column vector x_i, we will write the vector as $x_i = [x_i^0\,;\,\bar{x}_i]$ with x_i^0 being the first component and $\bar{x}_i$ being the vector consisting of the remaining components. Given square matrices P, Q, the notation $[P\,;\,Q]$ means that Q is appended to the last row of P; and $\mathrm{diag}(P, Q)$ denotes the block diagonal matrix with P, Q as its diagonal blocks. Throughout this paper, $\|\cdot\|$ denotes the matrix 2-norm or vector 2-norm, unless otherwise specified. For a given matrix M, we let $\lambda_{\max}(M)$ and $\lambda_{\min}(M)$ be the largest and smallest eigenvalues of M in magnitude, respectively. The condition number of a matrix M (not necessarily square) is the number $\kappa(M) = \sigma_{\max}(M)/\sigma_{\min}(M)$, where $\sigma_{\max}(M)$ and $\sigma_{\min}(M)$ are the largest and smallest singular values of M, respectively. We use $O(\mu)$ to denote a quantity that can be bounded by a constant times μ, and $\Theta(\mu)$ to denote that there exist two positive constants c_1, c_2 such that $c_1\mu \le \Theta(\mu) \le c_2\mu$.

We consider the following standard primal and dual SOCP problems:

$$\begin{array}{lll} \text{(P)} & \min & c_1^T x_1 + c_2^T x_2 + \cdots + c_N^T x_N \\ \text{s.t.} & & A_1 x_1 + A_2 x_2 + \cdots + A_N x_N \;=\; b \\ & & x_i \succeq 0, \\ \text{(D)} & \max & b^T y \\ \text{s.t.} & & A_i^T y + z_i \;=\; c_i, \quad i = 1, \ldots, N \\ & & z_i \succeq 0, \end{array} \tag{1}$$

where $A_i \in I\!R^{m\times n_i}$, $c_i, x_i, z_i \in I\!R^{n_i}$, $i = 1, \ldots, N$, and $y \in I\!R^m$. The constraint $x_i \succeq 0$ is a second order cone constraint defined by $x_i^0 \ge \|\bar{x}_i\|$. In particular, if the cone dimension n_i is 1, then the constraint $x_i \succeq 0$ is simply the standard non-negativity constraint $x_i \ge 0$, and such a variable is called a linear variable.

For convenience, we define

$$A \;=\; [A_1 \;\; A_2 \;\; \cdots \;\; A_N], \quad c \;=\; [c_1\,;\,c_2\,;\,\cdots\,;\,c_N],$$
$$x \;=\; [x_1\,;\,x_2\,;\,\cdots\,;\,x_N], \quad z \;=\; [z_1\,;\,z_2\,;\,\cdots\,;\,z_N], \quad n \;=\; \textstyle\sum_{i=1}^N n_i.$$

Throughout this paper, the notation $x \succeq 0$ ($x \succ 0$) is used to mean that each x_i is in (the interior of) the ith second order cone.

The perturbed KKT conditions of the primal-dual systems (1) are:

$$\begin{array}{rcll} Ax & = & b & \text{(primal feasibility)} \\ A^T y + z & = & c & \text{(dual feasibility)} \\ x_i \circ z_i & = & \nu e_i, \quad i = 1, \ldots, N, & \text{(perturbed complementary)} \\ x_i, z_i & \succeq & 0, & \end{array} \tag{2}$$

where e_i is the first unit vector in $\mathbb{R}^{n_i}$ and ν is a positive parameter that is to be driven to 0 explicitly. Here

$$x_i \circ z_i = [x_i^T z_i \ ; \ x_i^0 \bar{z}_i + z_i^0 \bar{x}_i].$$

In this paper, we will assume that A has full row rank, and that (P) and (D) in (1) are strictly feasible. Under these assumptions, as ν varies, the solutions to the perturbed KKT conditions (2) form a path (known as the central path) in the interior of the primal-dual feasible region, and as ν gradually reduces to 0, the path converges to an optimal solution of the primal and dual SOCP problems.

At each iteration of an IPM, the Newton equation associated with (2) needs to be solved. By performing block eliminations, one can either solve a system of linear equations of size $m+n$ or one of size m. These linear systems are known as the augmented and Schur complement equations (SCE), respectively. Currently, most implementations of IPMs [3, 22, 26] are based on solving the SCE since it has the obvious advantage of being smaller in size as well as being symmetric positive definite. However, as we shall see in section 3, the coefficient matrix in the SCE can be severely ill-conditioned when ν is close to 0, and this imposes a limit on how accurately one can solve an SOCP problem. In the case of LP, the ill-conditioning of the augmented equation was analyzed by Wright [29, 30]. Under certain assumptions including nondegeneracy, the computed search directions from the augmented equation in LP are shown to be sufficiently accurate for the IPM to converge to high accuracy. The structure of the ill-conditioning of the SCE arising from LP was analyzed in [14], and a stablization method based on performing Gaussian elimination with a certain pivoting order was proposed to transform the SCE into a better conditioned linear system of equations.

However, the ill-conditioning of the augmented equation and SCE in nonlinear conic programming seems to be much more complicated than that of LP. The potential numerical difficulties posed by the ill-conditioned SCE in SOCP were recognized by developers of solvers for SOCP such as [3, 4], [24], and [26]. It was also recognized by Goldfarb and Scheinberg [10] and that motivated them to propose and analyze a product-form Cholesky factorization for the Schur complement matrix (the sum of a sparse symmetric positive definite matrix and a possibly dense low rank matrix) in order to compute a stable Cholesky factor. Subsequently, Sturm [24] implemented the product-form Cholesky factorization in his code SEDUMI to solve the SCE arising at each iteration of a homogeneous self-dual (HSD) IPM. In addition, SEDUMI also employed sophisticated techniques to minimize numerical cancellations when computing the SCE and its factorization [24]. These sophisticated techniques typically greatly improve the stability of the SCE approach, but as we will see in Section 4, for certain

extreme cases, they do not entirely ameliorate the numerical difficulties resulting from the inherently ill-conditioned SCE.

The IPM code SEDUMI differs from standard infeasible interior-point methods in that it solves the homogeneous self-dual (HSD) embedding model. A natural question to ask is whether SEDUMI's unusually good performance arises from the inherent structure of the HSD model itself or from the sophisticated numerical techniques it uses in solving the SCE (or both). For a certain class of SOCPs with no strictly feasible primal/dual points, we show numerically in Section 4 that SEDUMI's superior performance can be explained by the structure of the HSD model itself. For some SOCPs with strictly feasible points, we shall also see in Section 4 that the sophisticated numerical techniques sometimes may offer only limited improvement in the attainable accuracy when compared to simpler technique used to solve the SCE.

Herein we propose a method to compute the search directions based on a reduced augmented equation that is derived by applying block row operations to the augmented equation, together with appropriate partitioning of the eigen-space of its (1,1) block. The reduced augmented equation is generally much smaller in size compared to the original augmented equation, and numerical experiments show that it can generally be solved reasonably efficiently by a judicious choice of symmetric indefinite system solvers together with a careful construction of the reduced augmented matrix by properly preserving the sparsity in the SOCP data. Our numerical results show that reduced augmented equation based IPMs are computationally more expensive (but not by a large factor) than SCE based IPMs. However, our reduced augmented equation based IPMs are superior to SCE based IPMs in that the former can usually compute approximate optimal solutions that are much more accurate than the latter before numerical difficulties are encountered. For example, for the `schedxxx` SOCP problems selected from the DIMACS library [18], our reduced augmented equation based IPMs are able to compute approximate optimal solutions with accuracies of 10^{-9} or better, while the SCE based IPMs (SDPT3 version 3.1 and SEDUMI) can only compute solutions with accuracies of 10^{-3} or 10^{-4} in some cases.

The paper is organized as follows. In the next section, we introduce the augmented and Schur complement equations from which the search direction in each IPM iteration may be computed. In section 3, we present an analysis of the conditioning and the growth in the norm of the Schur complement matrix and relate that to the deterioration of the primal infeasibility as the interior-point iterates approach optimality. In section 4, we present numerical results obtained from two different SCE-based primal-dual IPMs. In section 5, we derive a reduced augmented equation based on the augmented equation and a certain partitioning of the eigenvalues of its (1,1) block. In section 6, we analyze the conditioning of the reduced augmented matrix. In section 7, we discuss major computational issues for the efficient solution of the reduced augmented equation. Numerical results for an IPM with search directions computed from the reduced augmented equation are presented in section 8. We present our conclusion in section 9.

2 The augmented and Schur complement equations

In this section, we present the linear systems that need to be solved to compute the search direction in each IPM iteration.

For x_i, we define

$$\boldsymbol{aw}(x_i) \;=\; \begin{bmatrix} x_i^0 & x_i^1 & \cdots & x_i^{n_i} \\ x_i^1 & x_i^0 & & \\ \vdots & & \ddots & \\ x_i^{n_i} & & & x_i^0 \end{bmatrix} \;=\; \begin{bmatrix} x_i^0 & \bar{x}_i^T \\ \bar{x}_i & x_i^0 I \end{bmatrix}, \qquad \gamma(x_i) \;=\; \sqrt{(x_i^0)^2 - \|\bar{x}_i\|^2}. \tag{3}$$

For a given ν, the perturbed KKT conditions (2) in matrix form are:

$$\begin{aligned} Ax &= b \\ A^T y + z &= c \\ \boldsymbol{aw}(x)\,\boldsymbol{aw}(z)e^0 &= \nu e^0, \end{aligned} \tag{4}$$

where $e^0 = [e_1\,;\, e_2\,;\, \cdots\,;\, e_N]$. The matrix $\boldsymbol{aw}(x) = \mathrm{diag}(\boldsymbol{aw}(x_1), \cdots, \boldsymbol{aw}(x_N))$ is a block diagonal matrix with $\boldsymbol{aw}(x_1), \cdots, \boldsymbol{aw}(x_N)$ as its diagonal blocks. The matrix $\boldsymbol{aw}(z)$ is defined similarly.

For reasons of computational efficiency that we will explain later, in most IPM implementations for SOCP, a block diagonal scaling matrix is usually applied to the perturbed complementarity equation in (4). In this paper, we apply the Nesterov-Todd (NT) scaling matrix [26] to produce the following equation:

$$\boldsymbol{aw}(Fx)\,\boldsymbol{aw}(F^{-1}z)e^0 \;=\; \nu e^0, \tag{5}$$

where $F = \mathrm{diag}(F_1, \cdots, F_N)$ is chosen such that

$$Fx \;=\; F^{-1}z \;=:\; v. \tag{6}$$

For details on the conditions that F must satisfy and other scaling matrices, we refer the reader to [16].

Let

$$f_i \;=\; \begin{bmatrix} f_i^0 \\ \bar{f}_i \end{bmatrix} \;:=\; \frac{1}{\sqrt{2\big(\gamma(x_i)\gamma(z_i) + x_i^T z_i\big)}} \begin{bmatrix} \dfrac{1}{\omega_i} z_i^0 + \omega_i x_i^0 \\ \dfrac{1}{\omega_i}\bar{z}_i - \omega_i \bar{x}_i \end{bmatrix},$$

where $\omega_i = \sqrt{\gamma(z_i)/\gamma(x_i)}$. (Note that $\gamma(f_i) = 1$.) The precise form of F_i is given by

$$F_i \;=\; \omega_i \begin{bmatrix} f_i^0 & \bar{f}_i^T \\ \bar{f}_i & I + \dfrac{\bar{f}_i \bar{f}_i^T}{1 + f_i^0} \end{bmatrix}. \tag{7}$$

Let $\mu = x^T z/N$ be the normalized complementarity gap. The Newton equation associated with the perturbed KKT conditions (4) with NT scaling is given by

$$\begin{bmatrix} A & 0 & 0 \\ 0 & A^T & I \\ VF & 0 & VF^{-1} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \end{bmatrix} = \begin{bmatrix} r_p \\ r_d \\ r_c \end{bmatrix}, \tag{8}$$

where

$$V = \boldsymbol{aw}(v), \quad r_p = b - Ax, \quad r_d = c - z - A^T y, \quad r_c = \sigma\mu e^0 - Vv.$$

Note that we have chosen ν to be $\nu = \sigma\mu$ for some parameter $\sigma \in (0, 1)$.

The solution $(\Delta x, \Delta y, \Delta z)$ of the Newton equation (8) is referred to as the search direction. In each iteration of an IPM, solving (8) for the search direction is computationally the most expensive step. Observe that by eliminating Δz, the Newton equation (8) reduces to the so-called augmented equation:

$$\begin{bmatrix} -F^2 & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} r_x \\ r_p \end{bmatrix}, \tag{9}$$

where $r_x = r_d - FV^{-1}r_c$.

The augmented equation can further be reduced in size by eliminating Δx in (9) to produce the SCE:

$$\underbrace{AF^{-2}A^T}_{M} \Delta y \;=\; r_y \;:=\; r_p + AF^{-2}r_x \;=\; r_p + AF^{-2}r_d - AF^{-1}V^{-1}r_c. \tag{10}$$

The coefficient matrix $M = AF^{-2}A^T$ in (10) is known as the Schur complement matrix. It is symmetric and is positive definite as long as $x, z \succ 0$. The search direction corresponding to (8) always exists as long as $x, z \succ 0$. Note that if the scaling matrix F is not applied to the perturbed complementarity equation in (4), the corresponding Schur complement matrix would be $A\,\boldsymbol{aw}(z)^{-1}\boldsymbol{aw}(x)\,A^T$, which is a nonsymmetric matrix. It is well known that solving a linear system involving such a sparse nonsymmetric matrix is usually much more expensive than one involving a sparse symmetric positive definite matrix. Furthermore, the nonsymmetric coefficient matix is not guaranteed to be nonsingular even when $x, z \succ 0$. This explains why a suitable scaling matrix is usually applied to the perturbed complementarity equation in (4) before computing the search direction.

In its simplest form, most current implementations of IPMs compute the search direction $(\Delta x, \Delta y, \Delta z)$ based on the SCE (10) via the following procedure.

Simplified SCE approach:

(i) Compute the Schur complement matrix M and the vector r_y;

(ii) Compute the Cholesky or sparse Cholesky factor of M;
(iii) Compute Δy by solving 2 triangular linear systems involving the Cholesky factor;
(iv) Compute Δz from the equation $\Delta z = r_d - A^T \Delta y$;
(v) Compute Δx from the equation $\Delta x = F^{-2}(A^T \Delta y - r_x)$.

We should note that various heuristics to improve the numerical stability of the above simplified SCE approach are usually incorporated in the actual implementations. We will describe in section 4 variants of the simplified SCE approach implemented in two publicly available SOCP solvers, SDPT3, version 3.1 [27] and SEDUMI, version 1.05 [23].

The advantage of using the SCE is that it is usually a much smaller system compared to the augmented equation (9) or the Newton equation (8). Furthermore, the Schur complement matrix has the highly desirable property of being symmetric positive definite. In contrast, the coefficient matrix in (9) is symmetric but indefinite while that of (8) is nonsymmetric. The SCE is preferred also because it can be solved very efficiently via Cholesky or sparse Cholesky factorization of M, and there are highly efficient and machine optimized (yet user friendly) sparse Cholesky codes readily available, even in the public domain, the prime example being the sparse Cholesky codes of Ng and Peyton [17].

3 Conditioning of M and the deterioration of primal infeasibility

Despite the advantages of the SCE approach described in the last section, the SCE is however, generally severely ill-conditioned when the iterates (x, y, z) approach optimality, and this typically causes numerical difficulties. The most common numerical difficulty one may encounter in practice is that the Schur complement matrix M is numerically indefinite, although in exact arithmetic M is positive definite. Furthermore, the computed solution Δy from (10) may also be very inaccurate in that the residual norm $\|r_y - M\Delta y\|$ is much larger than the machine epsilon, and this typically causes the IPM to stall.

In this section, we will analyze the relationship between the norm $\|M\|$, the residual norm $\|r_y - M\Delta y\|$ of the computed solution Δy, and the primal infeasibility $\|r_p\|$, as the interior-point iterates approach optimality.

3.1 Eigenvalue decomposition of F^2

To analyze the norm $\|M\|$ and the conditioning of M, we need to know the eigenvalue decomposition of F^2.

Recall that $F = \mathrm{diag}(F_1, \cdots, F_N)$. Thus to find the eigenvalue decomposition of F^2, it suffices to find the eigenvalue decomposition of F_i^2, where F_i is the matrix in (7). By noting that for cones of dimensions $n_i \geq 2$, F_i^2 can be written in the form

$$F_i^2 \;=\; \omega_i^2 \left(I + 2(f_i f_i^T - e_i e_i^T) \right),$$

the eigenvalue decomposition of F_i^2 can readily be found. (The case where $n_i = 1$ is easy, and $F_i^2 = z_i/x_i$.) Without going through the algebraic details, the eigenvalue decomposition of F_i^2 is

$$F_i^2 \;=\; Q_i \Lambda_i Q_i^T, \tag{11}$$

where

$$\Lambda_i \;=\; \omega_i^2 \,\mathrm{diag}\Big((f_i^0 - \|\bar{f}_i\|)^2,\, (f_i^0 + \|\bar{f}_i\|)^2,\, 1, \cdots, 1\Big), \tag{12}$$

$$Q_i \;=\; \begin{bmatrix} -\dfrac{1}{\sqrt{2}} & +\dfrac{1}{\sqrt{2}} & 0 & \cdots & 0 \\ \dfrac{1}{\sqrt{2}}\, g_i & \dfrac{1}{\sqrt{2}}\, g_i & q_i^3 & \cdots & q_i^{n_i} \end{bmatrix}, \tag{13}$$

with

$$g_i \;:=\; [g_i^0\,;\, \bar{g}_i] \;=\; \bar{f}_i/\|\bar{f}_i\| \;\in\; I\!R^{n_i - 1}. \tag{14}$$

Notice that the first eigenvalue is the smallest and the second is the largest since $\gamma(f_i) = 1$. The set $\{q_i^3, \cdots, q_i^{n_i}\}$ is an orthonormal basis of the subspace $\{u \in I\!R^{n_i-1} \,:\, u^T g_i = 0\}$. To construct such an orthonormal basis, one may first construct the $(n_i - 1) \times (n_i - 1)$ Householder matrix H_i [11] associated with the vector g_i, then the last $n_i - 2$ columns of H_i is such an orthonormal basis. The precise form of H_i will be given later in section 6.

3.2 Analysis on $\|M\|$ and the conditioning of M

Here we will analyze how fast the norm $\|M\|$ and the condition number of M grow when $\mu \downarrow 0$, i.e., when the interior-point iterates approach an optimal solution (x^*, y^*, z^*). To simplify the analysis, we will analyze $\|M\|$ and the conditioning of M under the assumption that strict complementarity holds at the optimal solution. Unless otherwise stated, we assume that $n_i \geq 2$ in this subsection.

Referring to (2), strict complementarity [2] means for each ith pair of optimal primal and dual solutions, x_i^* and z_i^*, we have $\gamma(x_i^*) + \|z_i^*\|$ and $\gamma(z_i^*) + \|x_i^*\|$ both positive. In other words, (a) either $\gamma(x_i^*) = 0$ or $z_i^* = 0$, but not both; and (b) either $\gamma(z_i^*) = 0$ or $x_i^* = 0$, but not both. Thus, strict complementarity fails if for some i, either z_i^* is at the origin and x_i^* is on the boundary of the cone, or x_i^* is at the origin and z_i^* is on the boundary of the cone.

Under the assumption that strict complementarity holds at the optimal solution, we have the following three types of eigenvalue structures for F_i^2 when $x_i \circ z_i = \mu e_i$ and μ is small, following the classification in [10]. Note that $x_i^T z_i = \mu$.

Type 1 solution: $x_i^* \succ 0$, $z_i^* = 0$. In this case, $\gamma(x_i) = \Theta(1)$, $\gamma(z_i) = \Theta(\mu)$, and $\omega_i = \Theta(\sqrt{\mu})$. Also, $f_i^0, \|\bar{f}_i\| = \Theta(1)$, implying that all the eigenvalues of F_i^2 are $\Theta(\mu)$.

Type 2 solution: $x_i^* = 0$, $z_i^* \succ 0$. In this case, $\gamma(x_i) = \Theta(\mu)$, $\gamma(z_i) = \Theta(1)$, and $\omega_i = \Theta(1/\sqrt{\mu})$. Also, $f_i^0, \|\bar{f}_i\| = \Theta(1)$, implying that all the eigenvalues of F_i^2 are $\Theta(1/\mu)$.

Type 3 solution: $\gamma(x_i^*) = 0$, $\gamma(z_i^*) = 0$, $x_i^*, z_i^* \neq 0$. In this case, $\gamma(x_i), \gamma(z_i) = \Theta(\sqrt{\mu})$, and $\omega_i = \Theta(1)$. This implies that $f_i^0, \|\bar{f}_i\| = \Theta(1/\sqrt{\mu})$. Thus the largest eigenvalue of F_i^2 is $\Theta(1/\mu)$ and by the fact that $\gamma(f_i) = 1$, the smallest eigenvalue of F_i^2 is $\Theta(\mu)$. The rest of the eigenvalues are $\Theta(1)$.

Let D be the diagonal matrix consisting of the eigenvalues of F^2 sorted in ascending order. Then we have $F^2 = QDQ^T$, where the columns of Q are the sorted eigenvectors of F^2. Let D be partition into $D = \mathrm{diag}(D_1, D_2, D_3)$ such that $\mathrm{diag}(D_1)$ consists of all the small eigenvalues of F^2 of order $\Theta(\mu)$, and $\mathrm{diag}(D_2)$, $\mathrm{diag}(D_3)$ consist of the remaining eigenvalues of order $\Theta(1)$ and $\Theta(1/\mu)$, respectively. We also partition the matrix Q as $Q = [Q^{(1)}\ ,\ Q^{(2)}\ ,\ Q^{(3)}]$. Then $\widetilde{A} := AQ$ is partitioned as $\widetilde{A} = [\widetilde{A}_1\ ,\ \widetilde{A}_2\ ,\ \widetilde{A}_3] = [AQ^{(1)}\ ,\ AQ^{(2)}\ ,\ AQ^{(3)}]$. With the above partitions, we can express M as

$$M \ = \ \sum_{i=1}^{3} \widetilde{A}_i D_i^{-1} \widetilde{A}_i^T \ = \ \widetilde{A}_1 D_1^{-1} \widetilde{A}_1^T + O(1). \tag{15}$$

Lemma 3.1 *If there are Type 1 or Type 3 solutions, then $\|M\| = \Theta(1/\mu)\|A\|^2$. In addition, if $\widetilde{A}_1$ does not have full row rank, then $\|M^{-1}\| \geq \Theta(1)/\sigma_{\min}^2(A)$, where $\sigma_{\min}(A)$ is the smallest singular value of A.*

Proof. From (15), we have $\|M\| = \Theta(1/\mu)\|\widetilde{A}_1\|^2$. Since $\|\widetilde{A}_1\| \leq \|A\|\|Q^{(1)}\| = \|A\|$, the required result follows. If $\widetilde{A}_1$ does not have full row rank, then the component matrix $\widetilde{A}_1 D_1^{-1} \widetilde{A}_1^T$ of M is singular, and thus the smallest eigenvalue of M, $\lambda_{\min}(M)$, is of the order at most $\Theta(1)\sigma_{\min}^2(A)$, which implies that $\|M^{-1}\| = 1/\lambda_{\min}(M) \geq \Theta(1)/\sigma_{\min}^2(A)$. ∎

Remark 3.1 *(a) Lemma 3.1 implies that the growth in $\|M\|$ is caused by F^2 having small eigenvalues of the order $\Theta(\mu)$.*

(b) If F^2 has eigenvalues of the order $\Theta(\mu)$ and $\widetilde{A}_1$ does not have full row rank, then $\kappa(M) \geq \Theta(1/\mu)\kappa(A)^2$. On the other hand, if $\widetilde{A}_1$ has full row rank (which implies that the number of eigenvalues of F^2 of order $\Theta(\mu)$ is at least m), then $\kappa(M) = \Theta(1)\kappa(A)^2$.

(c) If there are only Type 2 solutions (thus $x^* = 0$ and $z \succ 0$), then $\|M\| = \Theta(\mu)\|A\|^2$. In this case, we also have $\kappa(M) = \Theta(1)\kappa(A)^2$.

Based on the results in [2], we have the following theorem concerning the rank of $\widetilde{A}_1$ and $\widetilde{A}_2$. We refer the reader to [2] for the definitions of primal and dual degeneracies.

Theorem 3.1 *Suppose that (x^*, y^*, z^*) satisfies strict complementarity. If the primal optimal solution x^* is primal nondegenerate, then $\left[\widetilde{A}_1\ ,\ \widetilde{A}_2\right]$ has full row rank when μ is small. If the dual optimal solution (y^*, z^*) is dual nondegenerate, then $\widetilde{A}_1$ has full column rank when μ is small.*

Proof. The result follows from Theorems 20 and 21 in [1]. ∎

Remark 3.2 *From Remark 3.1, we see that if $\widetilde{A}_1$ does not have full row rank or if (x^*, y^*, z^*) does not consist exclusively of Type 2 solutions, then M is ill-conditioned in that $\kappa(M) \geq \Theta(1/\mu)\kappa(A)^2$. Thus even if (x^*, y^*, z^*) satisfies strict complementarity, and is primal and dual nondegenerate, the associated Schur complement matrix M does not necessarily have bounded condition number when $\mu \downarrow 0$, unless $\widetilde{A}_1$ has full row rank. In fact, $\kappa(M)$ may grow at least proportional to $\Theta(1/\mu)\kappa(A)^2$. In contrast, for a linear programming problem (for which all the cones have dimensions $n_i = 1$), primal and dual nondegeneracy imply that $\widetilde{A}_1$ has both full column and row rank, and thus its associated Schur complement matrix M has bounded condition number when $\mu \downarrow 0$.*

3.3 Analysis of the deterioration of primal infeasibility

Although Cholesky factorization is stable for any symmetric positive definite matrix, the conditioning of the matrix may still affect the accuracy of the computed solution of the SCE. It is a common phenomenon that for SOCP, the accuracy of the computed search direction deteriorates as μ decreases due to an increasingly ill-conditioned M. As a result of this loss of accuracy in the computed solution, the primal infeasibility $\|r_p\|$ typically increases or stagnates when the IPM iterates approach optimality.

With the analysis of $\|M\|$ given in the last subsection, we will now analyze why the primal infeasibility may deteriorate or stagnate as interior-point iterations progress.

Lemma 3.2 *Suppose at the kth iteration, the residual vector in solving the SCE (10) is*

$$\xi \;=\; r_y - M\Delta y.$$

Assuming that Δx is computed exactly via the equation $\Delta x = F^{-2}(A^T\Delta y - r_x)$, then the primal infeasibility for the next iterate $x^+ = x + \alpha\Delta x$, $\alpha \in [0, 1]$, is given by

$$r_p^+ := b - Ax^+ \;=\; (1-\alpha)r_p \,+\, \alpha\xi.$$

Proof. We have

$$r_p^+ \;=\; (1-\alpha)r_p - \alpha(r_p - A\Delta x).$$

Now $A\Delta x = AF^{-2}(A^T\Delta y - r_x) = M\Delta y - r_y + r_p$, thus $r_p - A\Delta x = r_y - M\Delta y = \xi$ and the lemma is proved. ∎

Remark 3.3 *(a) In Lemma 3.2, we assume for simplicity that the component direction Δx is computed exactly. In finite precision arithmetic, errors will be introduced in the computation of Δx and that will also worsen the primal infeasibility r_p^+ of the next iterate in addition to $\|\xi\|$.*

(b) Observe that if the SCE is solved exactly, i.e., $\xi = 0$, then $\|r_p^+\| = (1-\alpha)\|r_p\|$, and the primal infeasibility should decrease monotoinically.

Lemma 3.2 implies that if the SCE is not solved to sufficient accuracy, then the inaccurate residual vector ξ may worsen the primal infeasibility of the next iterate. By standard perturbation error analysis, the worst-case residual norm of $\|\xi\|$ can be shown to be proportional to $\|M\|\|\Delta y\|$ times the machine epsilon u. The precise statement is given in the next lemma.

Lemma 3.3 *Let u be the machine epsilon. Given a symmetric positive definite matrix $B \in I\!R^{n\times n}$ with $(n+1)^2u \leq 1/3$, if Cholesky factorization is applied to B to solve the linear system $Bx = b$ and produced a computed solution $\widehat{x}$, then $(B+\Delta B)\widehat{x} = b$, for some ΔB with $\|\Delta B\|$ satisfying the following inequality:*

$$\|\Delta B\widehat{x}\| \leq 3(n+1)^2u\|B\|\|\widehat{x}\|.$$

Thus

$$\|b - B\widehat{x}\| \;=\; \|\Delta B\widehat{x}\| \;\leq\; O(n^2)u\|B\|\|\widehat{x}\|.$$

Proof. By Theorem 10.3 and 10.4, and their extensions in [13], $\|\Delta B\| \leq 2n\gamma_{n+1}(1 - n\gamma_{n+1})^{-1}\|B\|$, where $\gamma_{n+1} = \frac{(n+1)u}{1-(n+1)u}$. It is easy to show that $n\gamma_{n+1} \leq 1/3$ and that $2n\gamma_{n+1}(1-n\gamma_{n+1})^{-1} \leq 3(n+1)^2u$, and the required result follows. □

Remark 3.4 *Lemma 3.3 implies that if $\|B\|$ is large, then in the worst case scenario, the residual norm $\|b - B\widehat{x}\|$ is expected to be proportionately large.*

By Lemma 3.2 and the application of Lemma 3.3 to the SCE, we expect in the worst case that the primal infeasibility $\|r_p\|$ to grow at some level that is proportional to $\|M\|\|\Delta y\|u$. We end this section by presenting a numerical example to illustrate the relation between $\|r_p\|$ and $\|M\|\|\Delta y\|u$ in the last few iterations of an SCE based IPM when solving the random SOCP problem `rand200_800_1` (described in section 4).

The IPM we use is the primal-dual path-following method with Mehrotra predictor-corrector implemented in the MATLAB software SDPT3, version 3.1 [27]. But we should mention that to be consistent with the analysis presented in this section, the search directions are computed based on the simplifed SCE approach presented in section 2, not the more sophisticated variant implemented in SDPT3.

Table 1 shows the norms $\|M\|$, $\|M^{-1}\|$, $\|r_y - M\Delta y\|$ when solving the SCE (10). For this problem, $\|M\|$ and (hence $\kappa(M)$) grows like $\Theta(1/\mu)$ because its optimal solutions x_i^*, z_i^* are all of Type 3. The fifth and sixth columns in the table show that the residual norm in solving the SCE and $\|r_p\|$ deteriorate as $\|M\|$ increases. This is consistent with the conclusions of Lemmas 3.2 and 3.3. The last column further shows that $\|r_p\|$ increases proportionately to $\|M\|\|\Delta y\|u$, where the machine epsilon u is approximately 2.2×10^{-16}.

Figure 1 illustrates the phenomenon graphically for the SOCP problems `rand200_800_1` and `sched_50_50_orig`. The curves plotted correspond to the relative duality gap (`relgap`),

and the relative primal and dual infeasibility (`p-inf` and `d-inf`), defined by

$$\texttt{relgap} \;=\; \frac{|c^Tx - b^Ty|}{1 + (|c^Tx| + |b^Ty|)/2}, \quad \texttt{p-inf} \;=\; \frac{\|r_p\|}{1+\|b\|}, \quad \texttt{d-inf} \;=\; \frac{\|r_d\|}{1+\|c\|}. \tag{16}$$

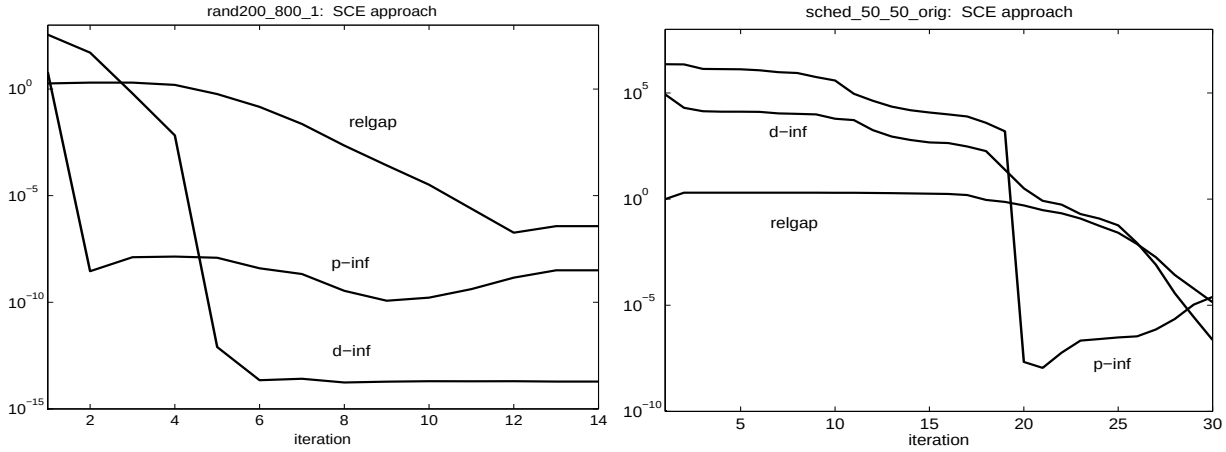

Figure 1: Convergence history of the SOCPs problems `rand200_800_1` and `sched_50_50_orig` when solved by the SCE based IPM in SDPT3, version 3.1. Notice that the relative primal infeasibility `p-inf` deteriorates as interior-point iterates approach optimality, while `relgap` may stagnate.

4 Computational results of two SCE based IPMs on solving some SOCP problems

Here we present numerical results for the SCE based IPMs implemented in the public domain solvers, SDPT3, version 3.1 [27], and SEDUMI, version 1.05 [23]. In this paper, all the numerical results are obtained from a Pentium IV 2.4GHz PC with 1G RAM running a Linux operating system.

Before we analyze the performance of the SCE based IPMs implemented in SDPT3 and SEDUMI, we must describe the methods employed to solve the SCE in both solvers. The IPM in SDPT3 is an infeasible path-following method that attempts to solve the central path equation based on (2), even if it does not exist. It solves the resulting SCE in each IPM iteration as follows. First it computes the Cholesky or sparse Cholesky factor of the Schur complement matrix M. Then the computed Cholesky factor is used to construct a preconditioner within a preconditioned symmetric quasi-minimal residual (PSQMR) Krylov subspace iterative solver employed to solve the SCE for Δy. The computations of Δz and Δx are the same as in the simplified SCE approach presented in section 2.

Table 1: The norm of the Schur complement matrix and $\|r_p\|$ associated with the last few IPM iterations for solving the SOCP problem `rand_200_800_1`.

Iter	$\|M\|$	$\|M^{-1}\|$	$\mu := x^T z/N$	$\|\Delta y\|$	$\|r_y - M\Delta y\|$	$\|r_p\|$	$\frac{\|r_p\|}{\|M\|\|\Delta y\|u}$
6	4.0e+11	2.7e+02	5.1e-04	2.2e+00	4.2e-05	9.3e-05	4.8e-01
7	3.4e+11	3.8e+02	8.1e-05	3.9e-01	7.1e-06	1.4e-05	4.8e-01
8	1.4e+12	5.6e+02	7.8e-06	8.4e-02	4.2e-06	5.0e-06	1.9e-01
9	1.8e+13	4.9e+02	9.2e-07	2.0e-02	8.7e-06	7.1e-06	8.9e-02
10	2.9e+14	2.3e+02	1.2e-07	2.7e-03	2.5e-05	1.8e-05	1.0e-01
11	3.8e+15	4.0e+01	1.1e-08	5.1e-04	4.9e-05	6.1e-05	1.4e-01
12	1.9e+17	6.8e+00	1.2e-09	9.4e-05	2.2e-04	1.4e-04	3.5e-02
13	1.2e+18	3.8e+01	1.8e-10	2.5e-03	3.7e-02	9.0e-04	1.4e-03

SEDUMI is a very well implemented SCE based public domain solver for both SOCP and SDP. The IPM in SEDUMI is not based on the central path for the original primal and dual problems (1), but that of the homogeneous self-dual (HSD) model of Ye, Todd, and Mizuno [31], by introducing 3 new auxiliary variables. The HSD model has the nice theoretical property that a strictly feasible primal and dual point always exists even if the original problems do not have one, and as a result the central path for the HSD model always exists, which is not necessarily true for the original problems in (1). As a consequence of this nice property, its solution set is always bounded. The same cannot be said for the original problems. For a problem that models an unrestricted variable by the difference of 2 nonnegative variables, the solution set for the original primal SOCP (P) is unbounded, and the feasible region of (D) has an empty interior, implying that the primal-dual central path does not exist. The HSD model, on the other hand, does not suffer from these defects. Thus the IPM in SEDUMI will not feel the effect of the unbounded solution set and nonexistence of the central path in the original problems in (1), but the effect of the unboundedness of the solution set on the infeasible path-following IPM in SDPT3 can be substantial and it often causes serious numerical difficulties.

The computation of the search direction in SEDUMI is based on the SCE associated with the HSD model. But it employs sophisticated numerical techniques to minimize numerical cancellations in its implementation of the SCE approach [24]. It computes the Schur complement matrix in the scaled space (called v-space) framework, and transform back and forth between quantities in the scaled and original spaces. It also employs a the product-form Cholesky (PFC) factorization [10] to factorize the Schur complement matrix. If the computed PFC factor is deemed sufficiently stable, SEDUMI will proceed to compute Δy by solving two triangular linear systems involving the PFC factor; otherwise, it will solve

the SCE by using the preconditioned conjugate gradient iterative method with a preconditioner constructed from the PFC factor. Note that the PFC factorization has been shown in [10] to produce stable triangular factors for the Schur complement matrix if the iterates are sufficiently close to the central path and strict complementarity holds at optimality. It is important to note, however, that using a stable method to solve the SCE does not necessarily imply that the computed direction $(\Delta x, \Delta y, \Delta z)$ based on the SCE approach will produce a small residual norm with respect to the original linear system (8); see Theorem 3.2 of [12] for the case of SDP.

We tested the SCE based IPMs in SDPT3 and SEDUMI on the following set of SOCP problems.

(a) The first set consists of 18 SOCPs collected by Pataki and Schmieta [18], which are available at `http://dimacs.rutgers.edu/Challenges/Seventh/Instances/`

(b) The second set consists of 10 SOCPs from the FIR Filter Optimization Toolbox of Scholnik and Coleman, available at `http://www.csee.umbc.edu/~dschol2/opt.html`

(c) The last set consists of 8 randomly generated SOCPs. The set of random problems `randxxx` are generated to be feasible and dominated by Type 3 solutions. For each problem, the constraint matrix A has the form $V_1 \Sigma V_2^T$, where V_1, V_2 are matrices whose columns are orthonormal, and Σ is a diagonal matrix with random diagonal elements drawn from the standard normal distribution, but a few of the diagonal elements are set to 10^5 to make A moderately ill-conditioned.

The statistics for the test problems are shown in Table 2. In the table, column "S" indicates the source of the problems: "D" stands for SOCP problems from the DIMACS library [18]. "F" stands for problems from the filter optimization toolbox [21]. "R" stands for randomly generated problems.

In our experiments, we stop the IPM iteration in SDPT3 when any of the following situations are encountered:

1. $\max(\texttt{relgap}, \texttt{p-inf}, \texttt{d-inf}) \leq 10^{-10}$;
2. incurable numerical difficulties (such as the Schur complement matrix being numerically indefinite) occurs;
3. `p-inf` has deteriorated to the extent that `p-inf` $>$ `relgap`.

SEDUMI also has a similar sets of stopping conditions but based on the variables of the HSD model. In SEDUMI, the dual conic constraints are not strictly enforced, thus the measure `d-inf` for SEDUMI is defined to be `d-inf` $= \max(\|r_d\|, \|z^-\|)$, where $\|z^-\|$ measures how much the dual conic constraints are violated.

We define

$$\phi \;:=\; \log_{10}(\max\{\texttt{relgap}, \texttt{p-inf}, \texttt{d-inf}\}). \tag{17}$$

Table 3 shows the numerical results for SDPT3 and SEDUMI on 36 SOCP problems. Observe that the accuracy exponent (ϕ) for many of the problems fall short of the targeted exponent of -10. For the `sched-xxx` problems, the accuracy exponents attained are especially poor, only -3 or -4 in some cases. We should mention that the results shown in Table 3 are not isolated to just the IPMs implemented in SDPT3 or SEDUMI, similar results are also observed in the SCE based IPM implemented by Andersen et al. [3]. For example, for the DIMACS SOCP problem `sched_50_50_orig`, the IPM in [3] reported the values 0.9 and 0.002 for the maximum violation of certain primal bound constraints and the dual constraints, respectively.

Table 2: Problem statistics of SOCP problems. An entry of the form "793x3" in the "SOC blk" column means that there are 793 3-dimensional second order cones. For the numbers in the "LIN blk" column, they indicate the number of linear variables, or 1-dimensional cones.

Problem	S	m	SOC blk	LIN blk
nb	D	123	793x3	4
nb_L1	D	915	793x3	797
nb_L2	D	123	1x1677, 838x3	4
nb_L2_bessel	D	123	1x123, 838x3	4
nql30	D	3680	900x3	3602
nql60	D	14560	3600x3	14402
nql180	D	130080	32400x3	129602
qssp30	D	3691	1891x4	2
qssp60	D	14581	7381x4	2
qssp180	D	130141	65341x4	2
sched_50_50_orig	D	2527	1x2474, 1x3	2502
sched_100_50_orig	D	4844	1x4741, 1x3	5002
sched_100_100_orig	D	8338	1x8235, 1x3	10002
sched_200_100_orig	D	18087	1x17884, 1x3	20002
sched_50_50_scaled	D	2526	1x2475	2502
sched_100_50_scaled	D	4843	1x4742	5002
sched_100_100_scaled	D	8337	1x8236	10002
sched_200_100_scaled	D	18086	1x17885	20002
firL1Linfalph	F	3074	5844x3	0
firL1Linfeps	F	7088	4644x3	1
firL1	F	6223	5922x3	0
firL2a	F	1002	1x1003	0
firL2L1alph	F	5868	1x3845, 1922x3	1
firL2L1eps	F	4124	1x203, 3922x3	0
firL2Linfalph	F	203	1x203, 2942x3	0
firL2Linfeps	F	6086	1x5885, 2942x3	0

Table 2: Problem statistics of SOCP problems. An entry of the form "793x3" in the "SOC blk" column means that there are 793 3-dimensional second order cones. For the numbers in the "LIN blk" column, they indicate the number of linear variables, or 1-dimensional cones.

Problem	S	m	SOC blk	LIN blk
firL2	F	102	1x103	0
firLinf	F	402	3962x3	0
rand200_300_1	R	200	20x15	0
rand200_300_2	R	200	20x15	0
rand200_800_1	R	200	20x40	0
rand200_800_2	R	200	20x40	0
rand400_800_1	R	400	40x20	0
rand400_800_2	R	400	40x20	0
rand700_1e3_1	R	700	70x15	0
rand700_1e3_2	R	700	70x15	0

Table 3: Accuracy attained by 2 SCE based IPMs for solving SOCP problems. The timings reported are in seconds. A number of the form "1.7-4" means 1.7×10^{-4}.

		SDPT3					SeDuMi			
problem	ϕ	`relgap`	`p-inf`	`d-inf`	Time	ϕ	`relgap`	`p-inf`	`d-inf`	Time
nb	-3.6	2.2-4	1.3-5	8.3 -9	8.3	-11.3	6.5-13	4.8-12	2.8-15	13.7
nb-L1	-5.0	5.8-6	1.0-5	9.6-11	14.3	-12.2	6.2-13	1.2-14	1.5-14	14.6
nb-L2	-5.5	5.1-8	3.1-6	8.9-12	12.7	-9.3	5.4-10	3.1-12	9.7-12	29.2
nb-L2-bessel	-6.4	3.1-7	4.3-7	8.8-11	7.3	-10.5	3.3-11	8.0-14	1.7-13	20.3
nql30	-6.3	4.9-7	7.8-8	2.2-12	5.0	-10.2	6.8-11	3.4-11	3.4-11	2.9
nql60	-6.5	3.4-7	1.3-7	2.8-12	22.2	-10.0	1.0-10	1.1-11	1.1-11	12.7
nql180	-5.3	4.9-6	8.5-7	5.7-12	265.0	-9.2	5.8-10	1.9-11	1.9-11	229.0
qssp30	-8.7	2.0 -9	2.3-10	2.2-14	5.0	-11.1	7.1-13	4.8-12	7.5-12	5.1
qssp60	-7.9	1.4-8	1.7 -9	3.4-15	25.7	-10.6	3.3-12	1.7-11	2.7-11	28.8
qssp180	-7.6	2.8-8	4.8-10	1.7-14	449.8	-11.2	7.0-12	1.2-12	1.8-12	666.5
sched-50-50-ori	-4.3	2.7-6	4.8-5	1.2-8	5.7	-7.0	1.0-12	1.0-7	4.0-14	7.4
sched-100-50-or	-3.8	3.6-5	1.6-4	9.9-10	13.0	-6.0	2.9-13	1.0-6	1.4-12	16.2
sched-100-100-o	-2.8	3.5-4	1.6-3	6.0-6	22.8	-3.3	6.6-11	4.6-4	4.1-11	33.0
sched-200-100-o	-5.2	5.5-6	5.8-6	1.5-8	82.1	-3.9	4.8-12	1.2-4	2.3-11	67.5
sched-50-50-sca	-6.2	3.7 -9	6.3-7	4.4-15	5.6	-8.2	1.0-13	7.0 -9	1.1-14	9.2
sched-100-50-sc	-6.1	8.8-7	4.3 -9	6.8-14	11.7	-8.9	1.1-11	1.3 -9	1.2-11	23.7
sched-100-100-s	-6.2	2.1-7	5.7-7	3.5-14	21.1	-7.1	3.2-12	7.8-8	5.3-15	36.3

Table 3: Accuracy attained by 2 SCE based IPMs for solving SOCP problems. The timings reported are in seconds. A number of the form "1.7-4" means 1.7×10^{-4}.

		SDPT3					SeDuMi			
problem	ϕ	relgap	p-inf	d-inf	Time	ϕ	relgap	p-inf	d-inf	Time
sched-200-100-s	-5.8	3.2-8	1.7-6	1.4-13	59.9	-7.8	1.2-12	1.7-8	3.8-14	117.2
firL1Linfalph	-9.9	7.1-11	1.3-10	1.0-15	252.5	-4.7	5.2-7	1.8-5	0.0-16	238.3
firL1Linfeps	-10.2	5.8-11	5.3-11	6.5-16	231.7	-10.4	3.4-13	4.3-11	1.2-14	92.9
firL1	-10.1	2.3-11	7.3-11	1.0-15	641.8	-9.0	1.8-11	1.0 -9	8.1-12	528.2
firL2a	-10.3	5.0-11	6.5-16	0.8-16	38.2	-12.6	7.8-15	2.7-13	4.5-15	22.5
firL2L1alph	-10.1	7.9-11	5.6-12	7.4-16	118.7	-3.3	6.2-7	5.1-4	4.4-11	177.8
firL2L1eps	-10.4	3.7-11	2.3-11	0.9-15	188.4	-9.3	1.3-11	4.8-10	1.1-11	174.8
firL2Linfalph	-10.1	7.9-11	7.6-11	7.8-16	144.5	-9.5	4.0-12	3.5-10	8.1-14	198.3
firL2Linfeps	-10.2	7.1-11	2.0-11	7.1-16	321.7	-9.1	8.1-10	5.7-10	0.0-16	232.1
firL2	-11.3	5.2-12	4.7-16	2.0-16	0.3	-13.1	1.3-15	7.3-14	3.3-15	0.2
firLinf	-8.9	1.1 -9	1.2 -9	1.0-15	497.2	-9.3	6.6-13	5.4-10	2.4-13	769.6
rand200-300-1	-8.0	1.1-8	1.3 -9	6.0-15	2.9	-6.4	3.3-7	4.3-7	0.0-16	9.9
rand200-300-2	-6.4	3.8-7	5.0 -9	5.9-14	2.7	-5.0	7.8-6	9.0-6	0.0-16	17.0
rand200-800-1	-4.5	3.3-5	2.6-10	2.0-14	5.0	-5.0	1.0-5	1.1-6	0.0-16	33.1
rand200-800-2	-4.1	7.4-5	3.7 -9	7.8-14	5.7	-5.8	8.8-7	1.6-6	0.0-16	73.7
rand400-800-1	-5.5	3.2-6	3.9 -9	6.2-10	16.9	-5.1	7.1-6	2.7-6	0.0-16	35.6
rand400-800-2	-5.3	5.0-6	1.5 -9	4.3 -9	17.2	-4.5	3.5-5	2.2-5	0.0-16	71.6
rand700-1e3-1	-5.9	1.2-6	4.7 -9	3.0-14	67.7	-5.7	1.6-6	1.9-6	0.0-16	174.9
rand700-1e3-2	-5.5	2.9-6	1.8-8	7.0-14	74.6	-4.6	1.4-5	2.6-5	0.0-16	250.7

From Table 3, we have thus seen the performance of SCE based IPMs for two rather different implementations in SDPT3 and SEDUMI. It is worthwhile to analyze the performance of these implementations to isolate the factor contributing to the good performance in one implementation, but not the other. On the first 10 SOCP problems, `nbxxx`, `nqlxxx` and `qsspxxx` in the DIMACS library, SEDUMI performs vastly better than the IPM in SDPT3 in terms of the accuracy attained by the approximate optimal solutions. We hypothesize that SEDUMI is able to obtain accurate approximate optimal solutions for these test problems primarily because of nice theoretical properties (existence of a strictly feasible point, and boundedness of solution set) of the HSD model. These problems contain linear variables that are the results of modeling unrestricted variables as the difference of two nonnegative vectors. Consequently, the resulting primal SOCPs have unbounded solution sets and the feasible regions of the dual SOCPs have empty interior. It should come as no surprise that the IPM in SDPT3 has trouble solving such a problem to high accuracy since the ill-conditioning in the Schur complement matrix is made worse by the growing norm of

the primal linear variables as the iterates approach optimality. On the other hand, for the IPM in SEDUMI, the ill-conditioning of the Schur complement matrix is not amplified since the norm of the primal variables in the HSD model stays bounded.

To verify the above hypothesis, we solve the `nbxxx`, `nqlxxx` and `qsspxxx` problems again in SDPT3, but in each IPM iteration, we trim the growth in the primal linear variables, x^u_+, x^u_-, arising from unrestricted variables x_u using the following heuristic [27]:

$$x^u_+ \;:=\; x^u_+ - 0.8\min(x^u_+, x^u_-), \qquad x^u_- \;:=\; x^u_- - 0.8\min(x^u_+, x^u_-). \tag{18}$$

This modification does not change the original variable x^u but it slows down the growth of x^u_+, x^u_-. After these modified vectors have been obtained, we also modify the associated dual linear variables z^u_+, z^u_- as follows if $\mu \le 10^{-4}$:

$$(z^u_+)_i \;:=\; \frac{0.5\mu}{\max(1,(x^u_+)_i)}, \qquad (z^u_-)_i \;:=\; \frac{0.5\mu}{\max(1,(x^u_-)_i)}. \tag{19}$$

Such a modification in z^u_+, z^u_- ensures that they approach 0 at the same rate as μ, and thus prevents the dual problem from attaining the equality constraints in (D) prematurely.

The results shown in Table 4 supported our hypothesis. Observe that with the heuristic in (18) and (19) to control the growth of $(x^u_+)_i/(z^u_+)_i$ and $(x^u_-)_i/(z^u_-)_i$, the IPM in SDPT3 can also achieve accurate approximate solutions, just like what the IPM based on the HSD model in SEDUMI is able to achieve. It is surprising that such a simple heuristic to control the growth can result in such a dramatic improvement on the achievable accuracy, even though the problems (P) and (D) in (1) do not have a strictly feasible point and the corresponding central path does not exist.

On other problems such as `schedxxx`, `firxxx`, and `randxxx`, the performance of SDPT3 and SEDUMI is quite comparable in terms of accuarcy attained, although SEDUMI is generally more accurate on the `schedxxx` problems, while SDPT3 performs somewhat better on the `randxxx` problems in terms of accuracy and computation time. On the `firxxx` problems, SDPT3 seems to be more robust whereas SEDUMI runs into numerical difficulties quite early when solving `firL1Linfalph` and `firL2L1alph`.

The empirical evidences produced by Table 3 show that the highly sophisticated numerical techniques employed to solve the SCE in SEDUMI can help to achieve better accuracy, but sometimes these techniques give limited improvement over the less sophisticated techniques employed to solve the SCE in SDPT3. On SOCP problems where the two solvers have vastly different performance in terms of accuracy, the difference can be attributed to the inherent IPM models used in the solvers rather than the numerical techniques employed to solve the SCE. The conclusion we may draw here is that the SCE is generally inherently ill-conditioned, and if our wish is to compute the search direction of (8) to higher accuracy, a new approach other than the SCE is necessary.

Table 4: Performance of the SCE based IPM in SDPT3 in solving SOCP problems with linear variables coming from unrestricted variables. The heuristics in (18) and (19) are applied at each IPM iteration.

		SDPT3					SeDuMi			
problem	ϕ	relgap	p-inf	d-inf	Time	ϕ	relgap	p-inf	d-inf	Time
nb-u	-10.2	6.4-11	1.1-13	5.2-16	16.7	-11.1	6.5-13	8.4-12	0.0-16	14.0
nb-L1-u	-10.0	9.9-11	1.1-11	2.2-16	31.9	-12.2	6.1-13	1.0-14	1.0-14	15.1
nb-L2-u	-10.2	5.8-11	1.6-11	6.6-16	19.7	-9.3	5.4-10	3.1-12	6.5-12	29.6
nb-L2-bessel-u	-10.2	6.7-11	3.3-11	3.3-16	15.1	-10.5	3.3-11	7.9-14	1.7-13	20.4
nql30-u	-10.1	8.5-11	1.6-13	8.0-13	7.6	-10.2	6.8-11	3.4-11	2.8-11	4.4
nql60-u	-10.4	4.2-11	6.1-12	2.8-13	32.0	-10.0	1.0-10	1.1-11	8.9-12	12.8
nql180-u	-8.0	1.0-8	1.7-11	1.5-11	373.6	-9.2	5.8-10	1.9-11	1.1-11	294.7
qssp30-u	-10.0	7.4-11	9.2-11	2.7-15	5.0	-11.3	7.1-13	4.8-12	5.2-12	4.9
qssp60-u	-8.8	1.4 -9	1.8 -9	4.0-14	23.1	-10.8	3.3-12	1.7-11	1.7-11	28.6
qssp180-u	-9.0	1.1 -9	6.7-10	9.5-15	539.9	-11.2	7.0-12	1.2-12	9.9-13	647.8

5 Reduced augmented equation

In this section, we present a new approach to compute the search direction via a potentially better-conditioned linear system of equations. Based on the new approach, the accuracy of the computed search direction is expected to be better than that computed from the SCE when μ is small.

In this approach, we start with the augmented equation in (9). By using the eigenvalue decomposition, $F^2 = QDQ^T$ presented in section 3.1, where $Q = \mathrm{diag}(Q_1, \cdots, Q_N)$ and $D = \mathrm{diag}(\Lambda_1, \cdots, \Lambda_N)$. We can diagonalize the (1,1) block and rewrite the augmented equation (9) as follows.

$$\begin{bmatrix} -D & \widetilde{A}^T \\ \widetilde{A} & 0 \end{bmatrix} \begin{bmatrix} \Delta\widetilde{x} \\ \Delta y \end{bmatrix} = \begin{bmatrix} \widetilde{r} \\ r_p \end{bmatrix}, \tag{20}$$

where

$$\widetilde{A} = AQ, \quad \Delta\widetilde{x} = Q^T\Delta x, \quad \widetilde{r} = Q^T r_x. \tag{21}$$

The augmented equation (20) has dimension $m+n$, which is usually much larger than m, the dimension of the SCE. We can try to reduce its size while overcoming some of the undesirable features of the SCE such as the growth of $\|M\|$ when $\mu \downarrow 0$.

Let the diagonal matrix D be partitioned into two parts as $D = \mathrm{diag}(D_1, \bar{D}_1)$ with $\mathrm{diag}(D_1)$ consisting of the small eigenvalues of F^2 of order $\Theta(\mu)$ and $\mathrm{diag}(\bar{D}_1)$ consisting of the remaining eigenvalues of magnitudes at least $\Theta(1)$. We partition the eigenvector matrix Q accordingly as $Q = [Q^{(1)} \ , \ \bar{Q}^{(1)}]$. Then $\widetilde{A}$ is partitioned as $\widetilde{A} = [\widetilde{A}_1 \ , \ \bar{A}_1] = [AQ^{(1)} \ , \ A\bar{Q}^{(1)}]$ and $\widetilde{r} = [\widetilde{r}_1 \, ; \, \bar{r}_1] = [(Q^{(1)})^T r_x \, ; \, (\bar{Q}^{(1)})^T r_x]$. Similarly, $\Delta\widetilde{x}$ is partitioned as $\Delta\widetilde{x} = [\Delta\widetilde{x}_1 \, ; \, \Delta\bar{x}_1] = [(Q^{(1)})^T \Delta x \, ; \, (\bar{Q}^{(1)})^T \Delta x]$.

By substituting the above partitions into (20), and eliminating $\Delta\bar{x}_1$, it is easy to show that solving the system (20) is equivalent to solving the following:

$$\begin{bmatrix} \bar{A}_1 \bar{D}_1^{-1} \bar{A}_1^T & \widetilde{A}_1 \\ \widetilde{A}_1^T & -D_1 \end{bmatrix} \begin{bmatrix} \Delta y \\ \Delta\widetilde{x}_1 \end{bmatrix} = \begin{bmatrix} r_p + \bar{A}_1 \bar{D}_1^{-1} \bar{r}_1 \\ \widetilde{r}_1 \end{bmatrix}, \tag{22}$$

$$\Delta\bar{x}_1 \ = \ \bar{D}_1^{-1}(\bar{A}_1^T \Delta y - \bar{r}_1). \tag{23}$$

By its construction, the coefficient matrix in (22) does not have large elements when $\mu \downarrow 0$. But its (1,1) block is generally singular or nearly singular, especially when μ is close to 0. Since a singular (1,1) block is not conducive for symmetric indefinite factorization of the matrix or the construction of preconditioners for the matrix, we will construct an equivalent system with a (1,1) block that is less likely to be singular. Let E_1 be a given positive definite diagonal matrix with the same dimension as D_1. Throughout this paper, we take $E_1 = I$. Let $S_1 = E_1 + D_1$. By adding $\widetilde{A}_1 S_1^{-1}$ times the second block equation in (22) to the first block equation, we get

$$\widetilde{A}\mathrm{diag}(S^{-1}, \bar{D}_1^{-1})\widetilde{A}^T \Delta y + \widetilde{A}_1 S_1^{-1} E_1 \Delta\widetilde{x}_1 \ = \ r_p + \widetilde{A}\mathrm{diag}(S^{-1}, \bar{D}_1^{-1})\widetilde{r}.$$

This, together with the second block equation in (22) but scaled by $S_1^{-1/2}$, we get the following equivalent system:

$$\underbrace{\begin{bmatrix} \widetilde{M} & \widetilde{A}_1 S_1^{-1/2} \\ S_1^{-1/2} \widetilde{A}_1^T & -D_1 E_1^{-1} \end{bmatrix}}_{\mathcal{B}} \begin{bmatrix} \Delta y \\ S_1^{-1/2} E_1 \Delta\widetilde{x}_1 \end{bmatrix} = \begin{bmatrix} q \\ S_1^{-1/2} \widetilde{r}_1 \end{bmatrix}, \tag{24}$$

$$\Delta\bar{x}_1 \ = \ \bar{D}_1^{-1}(\bar{Q}^{(1)})^T (A^T \Delta y - r_x), \tag{25}$$

where

$$\widetilde{M} \ = \ \widetilde{A}\,\mathrm{diag}(S_1^{-1}, \bar{D}_1^{-1})\widetilde{A}^T, \quad q \ = \ r_p + \widetilde{A}\mathrm{diag}(S_1^{-1}, \bar{D}_1^{-1})\widetilde{r}. \tag{26}$$

We call the system in (24) the **reduced augmented equation** (RAE). Note that once Δy and $\Delta\widetilde{x}_1$ are computed from (24) and $\Delta\bar{x}_1$ is computed from (25), Δx can be recovered through the equation $\Delta x = Q[\Delta\widetilde{x}_1 \, ; \, \Delta\bar{x}_1]$.

Remark 5.1 *(a) If the matrix D_1 is null, then the RAE (24) is reduced to the SCE (10).*

(b) $\mathcal{B}$ is a quasi-definite matrix [9, 28]. Such a matrix has the nice property that any symmetric reordering $\Pi\mathcal{B}\Pi^T$ has a "Cholesky factorization" $L\Lambda L^T$ where Λ is diagonal with both positive and negative diagonal elements.

Observe that the $(1,1)$ block, $\widetilde{M}$, in (24) has the same structure as the Schur complement matrix $M = \widetilde{A}\,\mathrm{diag}(D_1^{-1}, \bar{D}_1^{-1})\widetilde{A}^T$. But for $\widetilde{M}$, $\mathrm{diag}(S_1^{-1}, \bar{D}_1^{-1}) = O(1)$, whereas for M, $\mathrm{diag}(D_1^{-1}, \bar{D}_1^{-1}) = O(1/\mu)$. Because of this difference, the reduced augmented matrix $\mathcal{B}$ has bounded norm as $\mu \downarrow 0$, but $\|M\|$ is generally unbounded. Under certain conditions, $\mathcal{B}$ can be shown to have a condition number that is bounded independent of the normalized complementarity gap μ. The precise statements are given in the following theorems.

Theorem 5.1 *Suppose in (24) we use a partition such that* $\mathrm{diag}(D_1)$ *consist of all the eigenvalues of F^2 of the order $\Theta(\mu)$. If the optimal solution of (1) satisfies strict complementarity, then $\|\mathcal{B}\|$ satisfies the following inequality*

$$\|\mathcal{B}\| \;\le\; O(1)\,\|A\|^2.$$

Thus $\|\mathcal{B}\|$ is bounded independent of μ (as $\mu \downarrow 0$).

Proof. It is easy to see that

$$\|\mathcal{B}\| \;\;\le\;\; \sqrt{2}\,\max(\,\|\widetilde{M}\| \,+\, \|\widetilde{A}_1 S_1^{-1/2}\|,\, \|S_1^{-1/2}\widetilde{A}_1^T\| \,+\, \|D_1 E_1^{-1}\|\,).$$

Under the assumption that the optimal solution of (1) satisfies strict complementarity, then as $\mu \downarrow 0$, $\|D_1\| \downarrow 0$, and $\|\bar{D}_1^{-1}\| = O(1)$, so it is possible to find a constant $\tau \ge 1$ but of order $\Theta(1)$ such that:

$$\max(\,\|S_1^{-1}\|,\; \|\bar{D}_1^{-1}\|,\; \|D_1 E_1^{-1}\|\,) \;\le\; \tau.$$

Now

$$\|\widetilde{M}\| \;\;\le\;\; \|\widetilde{A}\|\max(\|S_1^{-1}\|, \|\bar{D}_1^{-1}\|)\|\widetilde{A}\| \;\le\; \tau\|\widetilde{A}\|^2$$

and

$$\begin{aligned}
\|S_1^{-1/2}\widetilde{A}_1^T\| \;\;&\le\;\; \|S_1^{-1/2}\|\|\widetilde{A}_1^T\| \;\le\; \tau\|\widetilde{A}_1\| \\
\|\widetilde{A}_1 S_1^{-1/2}\| \;\;&\le\;\; \|\widetilde{A}_1\|\|S_1^{-1/2}\| \;\le\; \tau\|\widetilde{A}_1\|,
\end{aligned}$$

thus we have

$$\|\mathcal{B}\| \;\le\; \tau\sqrt{2}\,\max(\,\|\widetilde{A}\|^2 + \|\widetilde{A}_1\|,\, \|\widetilde{A}_1\| + 1\,) \;\le\; \tau\sqrt{2}(\|A\| + 1)^2.$$

From here, the required result follows. ∎

Lemma 5.1 *The reduced augmented matrix $\mathcal{B}$ in (24) satisfies the following inequailty:*

$$\|\mathcal{B}^{-1}\| \;\le\; 2\sqrt{2}\max(\|\widetilde{M}^{-1}\|\,,\,\|W^{-1}\|),$$

where $W = B_1^T\widetilde{M}^{-1}B_1 + D_1E_1^{-1}$ with $B_1 = \widetilde{A}_1S_1^{-1/2}$.

Proof. From [20, p. 389], it can be deduced that

$$\mathcal{B}^{-1} \;=\; \begin{bmatrix} \widetilde{M}^{-1/2}(I-P)\widetilde{M}^{-1/2} & \widetilde{M}^{-1}B_1W^{-1} \\ W^{-1}B_1^T\widetilde{M}^{-1} & -W^{-1} \end{bmatrix},$$

where $P = \widetilde{M}^{-1/2}B_1W^{-1}B_1^T\widetilde{M}^{-1/2}$. Note that P satisfies the condition $0 \preceq P \preceq I$, i.e., P and $I-P$ are positive definite. By the definition of W, we have $0 \preceq W^{-1/2}B_1^T\widetilde{M}^{-1}B_1W^{-1/2} \preceq I$, and thus $\|\widetilde{M}^{-1/2}B_1W^{-1/2}\| \le 1$. This implies that

$$\begin{aligned} \|\widetilde{M}^{-1}B_1W^{-1}\| \;&\le\; \|\widetilde{M}^{-1/2}\|\,\|\widetilde{M}^{-1/2}B_1W^{-1/2}\|\,\|W^{-1/2}\| \;\le\; \|\widetilde{M}^{-1/2}\|\|W^{-1/2}\| \\ &\le\; \max\left(\|\widetilde{M}^{-1}\|\,,\,\|W^{-1}\|\right). \end{aligned}$$

It is easy to see that

$$\begin{aligned} \|\mathcal{B}^{-1}\| \;&\le\; \sqrt{2}\max\left(\|\widetilde{M}^{-1/2}(I-P)\widetilde{M}^{-1/2}\| + \|\widetilde{M}^{-1}B_1W^{-1}\|\,,\,\|W^{-1}B_1^T\widetilde{M}^{-1}\| + \|W^{-1}\|\right) \\ &\le\; 2\sqrt{2}\max(\|\widetilde{M}^{-1}\|\,,\,\|W^{-1}\|). \end{aligned}$$

□

Theorem 5.2 *Suppose in (24) we use a partition such that* $\mathrm{diag}(D_1)$ *consist of all the eigenvalues of F^2 of the order $\Theta(\mu)$. If the optimal solution of (1) satisfies the strict complementarity, and the primal and dual nondegeneracy conditions defined in [2], then the condition number of the coefficient matrix in (24) is bounded independent of μ (as $\mu \downarrow 0$).*

Proof. Let $\bar{D}_1$ be further be partition into $\bar{D}_1 = \mathrm{diag}(D_2, D_3)$ where $\mathrm{diag}(D_2)$ and $\mathrm{diag}(D_3)$ consist of eigenvalues of F^2 of order $\Theta(1)$ and $\Theta(1/\mu)$, respectively. Let $\bar{Q}_1$ and $\bar{A}_1$ be partitioned accordingly as $\bar{Q}_1 = [Q_2\ ,\ Q_3]$ and $\bar{A}_1 = [AQ_2\ ,\ AQ_3] =: [\widetilde{A}_2\ ,\ \widetilde{A}_3]$. By Theorems 20 and 21 in [1], dual nondegeneracy implies that $\widetilde{A}_1 = AQ_1$ has full column rank and primal nondegeneracy implies that $[\widetilde{A}_1\ ,\ \widetilde{A}_2]$ has full row rank. Since $\widetilde{M} = [\widetilde{A}_1\,,\,\widetilde{A}_2]\mathrm{diag}(S_1^{-1},\,D_2^{-1})[\widetilde{A}_1\,,\,\widetilde{A}_2]^T + \Theta(\mu)$. Thus $\sigma_{\min}(\widetilde{M})$ is bounded away from 0 even when $\mu \downarrow 0$. This, together with the fact that $\widetilde{A}_1S_1^{-1/2}$ has full column rank, implies that the matrix $W := S_1^{-1/2}\widetilde{A}_1^T\widetilde{M}^{-1}\widetilde{A}_1S_1^{-1/2} + D_1E_1^{-1}$ has $\sigma_{\min}(W)$ bounded away from 0 even when $\mu \downarrow 0$. By Lemma 5.1, $\|\mathcal{B}^{-1}\|$ is bounded independent of μ. By Theorem 5.1, $\|\mathcal{B}\|$ is also bounded independent of μ, and the required result follows. □

6 Reduced augmented equation and primal infeasibility

Let $[\xi\,;\,\eta]$ be the residual vector for the computed solution of (24).

Lemma 6.1 *Let u be the machine epsilon and l be the dimension of $\Delta\widetilde{x}_1$. Suppose $(l+m)u \leq 1/2$ and we use Gaussian elimination with partial pivoting (GEPP) to solve (24) to get the computed solution $(\Delta\widetilde{x}_1\,;\,\Delta y)$, then the residual vector $[\xi\,;\,\eta]$ for the computed solution satisfies the following inequality:*

$$\|(\xi\,;\,\eta)\|_\infty \;\leq 4(l+m)^3 u\rho\,\|\mathcal{B}\|_\infty\|(\Delta\widetilde{x}_1\,;\,\Delta y)\|_\infty$$

where ρ is the growth factor associated with GEPP.

Proof. This lemma follows from Theorem 9.5 in [13]. ∎

Remark 6.1 *Theorem 5.1 stated that if strict complementarity holds at the optimal solution, then $\|\mathcal{B}\|_\infty$ will not grow as $\mu \downarrow 0$ in contrast to $\|M\|$, which usually grows proportionately to $\Theta(1/\mu)$. Now because the growth factor ρ for GEPP is usually $O(1)$, Lemma 6.1 implies that the residual norm $\|(\xi\,;\,\eta)\|_\infty$ will be maintained at some level proportional to $u\|A\|^2$ even when $\mu \downarrow 0$.*

Now we establish the relationship between the residual norm in solving (24) and the primal infeasibility associated with the search dierction computed from the RAE approach. Suppose that in computing $\Delta\bar{x}_1$ from (25), a residual vector δ is introduced, i.e.,

$$\Delta\bar{x}_1 \;\;=\;\; \bar{D}_1^{-1}(\bar{Q}^{(1)})^T(A^T\Delta y - r) - \delta.$$

Then we have the following lemma for the primal infeasibility of the next iterate.

Lemma 6.2 *Suppose Δx is computed from the RAE approach. Then the primal infeasibility $\|r_p^+\|$ for the next iterate $x^+ = x + \alpha\Delta x$, $\alpha \in [0,1]$, satisfies the following inequality:*

$$\|r_p^+\| \;\leq\; (1-\alpha)\|r_p\| + \alpha\|\xi + \bar{A}_1\delta - \widetilde{A}_1 S_1^{-1/2}\eta\|.$$

Proof. The computed solution from (24) satisfies

$$\widetilde{M}\Delta y + \widetilde{A}_1 S_1^{-1}E_1\Delta\widetilde{x}_1 \;=\; q - \xi \tag{27}$$

$$\widetilde{A}_1^T\Delta y - D_1\Delta\widetilde{x}_1 \;=\; \widetilde{r}_1 - S_1^{1/2}\eta. \tag{28}$$

From (27), we have

$$\bar{A}_1\bar{D}_1^{-1}\bar{A}_1^T\Delta y \;=\; q - \xi - S_1^{-1}E_1\Delta\widetilde{x}_1 - \widetilde{A}_1 S_1^{-1}\widetilde{A}_1^T\Delta y.$$

Thus we have

$$\begin{aligned} A\Delta x &= \widetilde{A}_1\Delta\widetilde{x}_1 + \bar{A}_1\bar{D}_1^{-1}(\bar{A}_1^T\Delta y - \bar{r}_1) - \bar{A}_1\delta \\ &= \widetilde{A}_1\Delta\widetilde{x}_1 - \bar{A}_1\bar{D}_1^{-1}\bar{r}_1 - \bar{A}_1\delta + q - \xi - \widetilde{A}_1S_1^{-1}E_1\Delta\widetilde{x}_1 - \widetilde{A}_1S_1^{-1}\widetilde{A}_1^T\Delta y. \end{aligned} \quad (29)$$

Now by (28), we have

$$\widetilde{A}_1S_1^{-1}\widetilde{A}_1^T\Delta y \;=\; \widetilde{A}_1S_1^{-1}(\widetilde{r}_1 - S_1^{1/2}\eta) + \widetilde{A}_1S_1^{-1}D_1\Delta\widetilde{x}_1.$$

By substituting the above equation into (29), we get

$$A\Delta x \;=\; r_p - \bar{A}_1\delta + \widetilde{A}_1S_1^{-1/2}\eta - \xi,$$

and the required result follows readily. □

Remark 6.2 *From Lemma 6.2, we see that if the RAE returns a small residual norm, then the primal infeasibility of the next iterate would not be seriously worsened by the residual norm. From Theorem 5.1 and Lemma 6.1, we expect the residual norm* $\|[\xi;\eta]\|$ *to be small since* $\|\mathcal{B}\|$ *is bounded independent of* μ*. Also, since by its construction,* $\bar{D}^{-1}$ *does not have large elements,* $\|\delta\|$ *is expected to be small as well.*

Figure 2 shows the convergence behavior of the IPM in SDPT3, but with search directions computed from the RAE (24) for the problems `ran200_800_1` and `sched_50_50_orig`. As can be seen from the relative primal infeasibility curves, the RAE approach is more stable than the SCE approach. It is worth noting that under the new approach, the solver is able to deliver 10 digits of accuracy, i.e. $\phi \leq -10$. This is significantly better than the accuray $\phi \approx -6$ attained by the SCE approach. Note that we use a partition such that eigenvalues of F^2 that are smaller than 10^{-3} are put in D_1.

In Table 5, we show the norms $\|\mathcal{B}\|$, $\|\mathcal{B}^{-1}\|$ and the residual norm in solving the RAE (24) for the last few IPM iterations in solving the problem `rand200_800_1`. The reader can notice that $\|\mathcal{B}\|$ and $\kappa(\mathcal{B})$ do not grow when $\mu \downarrow 0$ in contrast to $\|M\|$ and $\kappa(M)$ in Table 1. The residual norm for the computed solution of (24) remains small throughout, and in accordance with Lemma 6.1, the residual norm is approximately equal to $u\|\mathcal{B}\|$ times the norm of the computed solution. By Lemma 6.2, the small residual norm in solving the RAE explains why the primal infeasibility computed from the RAE approach does not deteriorate as in the SCE approach.

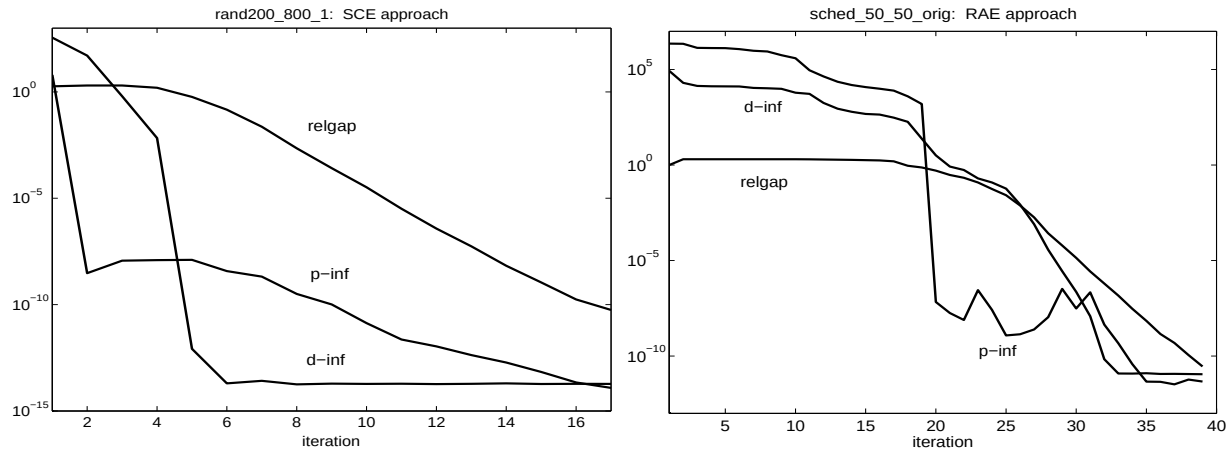

Figure 2: Same as Figure 1 but for the RAE approach in computing the search directions for the SOCP problems `rand200_800_1` and `sched_50_50_orig`. Notice that the primal infeasibility does not deteriorate when the iterates approach optimality. Both problems are primal and dual non-degenerate, and strict complementarity holds at optimality.

Table 5: Condition number of the reduced augmented matrix $\mathcal{B}$ associated with the last few IPM iterations for solving the SOCP problem `rand200_800_1`. The maximum number of columns in $\widetilde{A}_1$ is 19.

Iter	$\|\mathcal{B}\|$	$\|\mathcal{B}^{-1}\|$	$x^T z/N$	$\|[\Delta y; \Delta\widetilde{x}_1]\|$	residual norm	$\|r_p\|$	$\frac{\|r_p\|}{\|\mathcal{B}\|\|[\Delta y; \Delta\widetilde{x}_1]\|u}$
11	3.2e+11	3.1e+02	1.1e-08	1.3e-03	2.5e-08	4.6e-08	5.2e-01
12	3.7e+11	2.7e+02	1.3e-09	4.6e-04	6.9e-09	1.8e-08	4.7e-01
13	3.4e+11	2.8e+02	1.9e-10	2.3e-04	4.3e-09	8.1e-09	4.7e-01
14	2.6e+11	3.7e+02	2.5e-11	1.0e-04	1.2e-09	2.9e-09	4.9e-01
15	2.3e+11	4.2e+02	4.0e-12	3.6e-05	3.4e-10	9.4e-10	5.1e-01
16	2.0e+11	5.2e+02	5.6e-13	1.3e-05	1.4e-10	4.8e-10	8.6e-01

7 Major computational issues

The theoretical analysis in the last section indicates that the RAE approach is potentially more stable than the standard SCE approach, but the trade-off is that the former requires a larger indefinite linear system to be solved. Therefore, the efficiency in solving (24) is one of our major concern in the implementation.

In forming the reduced augmented matrix $\mathcal{B}$, those operations involving Q (the eigenvector matrix of F^2) must be handled carefully by exploiting the structure of Q to avoid incurring significant storage and computational cost. Also, the sparsity of AA^T must be properly preserved when computing $\widetilde{M}$.

7.1 Computations involving Q

The operations involving Q in assembling the RAE (24) are as follows:

- Computation of the (1,1) block $\widetilde{M} = AQ\mathrm{diag}(S_1^{-1}, \bar{D}_1^{-1})Q^T A^T$;
- Computation of the (1,2) block $\widetilde{A}_1 = AQ^{(1)}$;

the Householder matrix H_i explicitly. Without going into the algebraic details, the precise form of H_i is given as follows:

$$H_i \;=\; I - h_i h_i^T,$$

where

$$h_i \;=\; \frac{1}{\tau_i}\begin{bmatrix} \tau_i^2\,\mathrm{sign}(g_i^0) \\ \bar{g}_i \end{bmatrix} \;\in\; I\!R^{n_i-1}, \quad \tau_i := \sqrt{1+|g_i^0|}. \tag{30}$$

With some algebraic manipulations, the eigenvector Q_i can be rewritten in the form given in the next lemma. The explicit formula derived below is critical to the efficient computation of those operations involving Q mentioned above.

Lemma 7.1

$$Q_i \;=\; \mathrm{diag}(K_i, I) - u_i v_i^T, \tag{31}$$

where

$$K_i \;=\; \begin{bmatrix} -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \beta_i & \beta_i \end{bmatrix}, \quad u_i \;=\; \begin{bmatrix} 0 \\ h_i \end{bmatrix}, \quad v_i \;=\; \begin{bmatrix} \beta_i h_i^0 \\ \beta_i h_i^0 \\ \bar{h}_i \end{bmatrix}, \tag{32}$$

with $\beta_i = -\mathrm{sign}(h_i^0)/\sqrt{2}$.

Proof. Note that by construction, the first column of H_i is given by $-\mathrm{sign}(g_i^0)g_i$. Thus from (13), we have, by letting $\alpha = \frac{1}{\sqrt{2}} + \mathrm{sign}(g_i^0)$, that

$$\begin{aligned} Q_i \;&=\; \begin{bmatrix} -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 & \cdots & 0 \\ \frac{1}{\sqrt{2}}g_i & \alpha g_i & 0 & \cdots & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & I - h_i h_i^T \end{bmatrix} \\ &=\; \begin{bmatrix} -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 & \dots & 0 \\ \frac{g_i^0}{\sqrt{2}} & \alpha g_i^0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} + \begin{bmatrix} 0 \\ \bar{h}_i \end{bmatrix}\begin{bmatrix} \frac{\tau_i}{\sqrt{2}} \\ \alpha\tau_i \\ 0 \\ \vdots \\ 0 \end{bmatrix}^T - \begin{bmatrix} 0 \\ h_i \end{bmatrix}\begin{bmatrix} 0 \\ h_i \end{bmatrix}^T \\ &=\; \begin{bmatrix} -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 & \dots & 0 \\ -\frac{\mathrm{sign}(g_i^0)}{\sqrt{2}} & -\frac{\mathrm{sign}(g_i^0)}{\sqrt{2}} - 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} - \begin{bmatrix} 0 \\ h_i \end{bmatrix}\begin{bmatrix} -\frac{\tau_i}{\sqrt{2}} \\ h_i^0 - \alpha\tau_i \\ \bar{h}_i \end{bmatrix}^T. \end{aligned}$$

It is readily shown that $h_i^0 - \alpha\tau_i = -\tau_i/\sqrt{2} = \beta_i h_i^0$. Now it is easy to see that (31) and (32) hold. □

Observe that each Q_i is a rank one perturbation of a highly sparse block diagonal matrix. Based on the above lemma, those operations listed at the beginning of this subsection, except the first one, can be computed straightforwardly. To compute the matrix $\widetilde{M}$, we have to further analyze the structure of the matrix $Q_i \mathrm{diag}(S_{1i}^{-1}, \bar{D}_{1i}^{-1}) Q_i^T$.

Let $G_i = \mathrm{diag}(S_{1i}^{-1}, \bar{D}_{1i}^{-1})$ and $\Sigma_i = \mathrm{diag}(K_i, I_i)$, then $Q_i = \Sigma_i - u_i v_i^T$ and

$$Q_i\, G_i\, Q_i^T \;=\; \Sigma_i\, G_i\, \Sigma_i^T - \Sigma_i G_i v_i u_i^T - u_i v_i^T G_i \Sigma_i^T + u_i v_i^T\, G_i\, v_i u_i^T.$$

By setting $\widetilde{v}_i = \Sigma_i G_i v_i$, and $\rho_i = v_i^T G_i v_i$ (note that ρ_i is a positive scalar and Σ_i is an orthogonal matrix), we have

$$\begin{aligned} Q_i\, G_i\, Q_i^T &= \Sigma_i G_i \Sigma_i^T - \widetilde{v}_i u_i^T - u_i \widetilde{v}_i^T + \rho_i u_i u_i^T, \\ &= \Sigma_i G_i \Sigma_i^T + l_i l_i^T - \frac{1}{\rho_i} \widetilde{v}_i \widetilde{v}_i^T, \end{aligned} \tag{33}$$

where $l_i = \widetilde{v}_i/\sqrt{\rho_i} - \sqrt{\rho_i} u_i$. With the expression in (33), each component matrix $\widetilde{M}_i$ in $\widetilde{M} = \sum_{i=1}^N \widetilde{M}_i$ can be expressed as:

$$\widetilde{M}_i \;=\; A_i Q_i G_i Q_i^T A_i^T \;=\; A_i\, (\Sigma_i G_i \Sigma_i^T)\, A_i^T + (A_i l_i)(A_i l_i)^T - \frac{1}{\rho_i}(A_i \widetilde{v}_i)(A_i \widetilde{v}_i)^T. \tag{34}$$

Since $\Sigma_i G_i \Sigma_i^T$ is a highly sparse block diagonal matrix, thus $\widetilde{M}_i$ is a symmetric rank two perturbation to a sparse matrix if $A_i A_i^T$ is sparse. Hence, the computational complexity of $\widetilde{M}$ is only slightly more expensive than that for the Schur complement matrix M.

7.2 Handling dense columns

Let $\Sigma \;=\; \mathrm{diag}(\Sigma_1, \ldots, \Sigma_N)$, and

$$A_l \;=\; [A_1 l_1 \;,\; \ldots \;,\; A_N l_N], \quad A_v \;=\; \left[\frac{1}{\sqrt{\rho_1}} A_1 \widetilde{v}_1 \;,\; \ldots \;,\; \frac{1}{\sqrt{\rho_N}} A_N \widetilde{v}_N\right]. \tag{35}$$

Then it is readily shown that

$$\widetilde{M} \;=\; A\, (\Sigma \mathrm{diag}(S_1^{-1}, \bar{D}_1^{-1}) \Sigma^T)\, A^T + A_l A_l^T - A_v A_v^T. \tag{36}$$

If AA^T is sparse, then the first matrix in (36) is sparse as well. For an SOCP problem where all the cones are low dimensional, typically the matrices A_l and A_v are also sparse. In that case, the reduced augmented equation (24) may be solved directly. However, if high dimensional cones exist, then A_l and A_v invariably contain dense columns. Moreover, when

A is sparse but with dense columns, AA^T will also be dense. In order to preserve the sparsity in $\widetilde{M}$, it is necessary to handle the dense columns separately when they exist.

Let P_1 be the dense columns in $A\Sigma\mathrm{diag}(S_1^{-1/2}, \bar{D}_1^{-1/2})$ and A_l, and P_2 be the dense columns in A_v. Let $\widetilde{M}_s = \widetilde{M} - P_1P_1^T + P_2P_2^T$ be the "sparse part" of $\widetilde{M}$. It is well known that by introducing the following new variables

$$t_1 \;=\; P_1^T\Delta y, \qquad t_2 \;=\; -P_2^T\Delta y,$$

the dense columns can be removed from $\widetilde{M}$; see [4]. The precise form of the reduced augmented equation (24) with dense column handling is as follows:

$$\begin{bmatrix} \widetilde{M}_s & U \\ U^T & -C \end{bmatrix} \begin{bmatrix} \Delta y \\ S_1^{-1/2}E_1\Delta\widetilde{x}_1 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} q \\ S_1^{-1/2}\widetilde{r}_1 \\ 0 \\ 0 \end{bmatrix}, \tag{37}$$

where q is defined in (26) and

$$U = \left[\widetilde{A}_1S_1^{-1/2},\;\; P_1,\;\; P_2\right], \quad C = \mathrm{diag}(D_1E_1^{-1}, I_1, -I_2).$$

7.3 Direct solvers for symmetric indefinite systems

The solution of the sparse symmetric indefinite system (37) is one of the most computationally intensive step in each IPM iteration. Therefore, it is critical that the solver used must be as efficient as possible.

We consider two methods for solving (37). The first is the Schur complement method, which is also equivalent to the Sherman-Morrison-Woodbury formula. The second is the Bunch-Parlett factorization implemented in MA47 [19]. Each of these methods has its own advantage in different circumstances.

Schur complement method This method is widely used for dense column handling in IPM implementations; see [4] and the references therein. The method uses the sparse matrix $\widetilde{M}_s$ as the pivoting matrix to perform block eliminations in (37). It is readily shown that solving (37) is equivalent to solving the following systems:

$$\left(U^T\widetilde{M}_s^{-1}U + C\right)\begin{bmatrix} S_1^{-1/2}E_1\Delta\widetilde{x}_1 \\ t_1 \\ t_2 \end{bmatrix} \;=\; U^T\widetilde{M}_s^{-1}q - \begin{bmatrix} S_1^{-1/2}\widetilde{r}_1 \\ 0 \\ 0 \end{bmatrix} \tag{38}$$

$$\Delta y \;=\; \widetilde{M}_s^{-1}q - \widetilde{M}_s^{-1}U\begin{bmatrix} S_1^{-1/2}E_1\Delta\widetilde{x}_1 \\ t_1 \\ t_2 \end{bmatrix}.$$

Note that since $\widetilde{M}$ is symmetric positive definite, its "sparse part", $\widetilde{M}_s$, is typically also positive definite if the number of dense columns removed from $\widetilde{M}$ is small. If $\widetilde{M}_s$ is indeed positive definite, then (38) can be solved by Cholesky or sparse Cholesky factorization. As mentioned before, highly efficient and optimized sparse Cholesky solvers are readily available. Another advantage of the Schur complement method is that the symbolic factorization of $\widetilde{M}_s$ needs only be computed once or twice during the initial phase of the IPM iteration and it can be re-use for subsequent IPM iterations even when the partition used in D changes.

But the Schur complement method does have a major disadvantage in that the matrix $U^T\widetilde{M}_s^{-1}U + C$ is typically dense. This can lead to huge computational burden when U has large number of columns, say, more than a few hundreds. Furthermore, the Schur complement method is numerically less stable than a method that solves (37) directly.

Roughly speaking, the Schur complement method is best suited for problems with U having a small number of columns. When U has a large number of columns or when $\widetilde{M}_s$ is not positive definite, we have to solve (37) directly by the second method described below.

MA47 MA47 is a direct solver developed by J. Reid and I.S. Duff [19] for sparse symmetric indefinite systems. This is perhaps the only publicly available state of the art direct solver for sparse symmetric indefinite systems. It appears not to be as efficient as the sparse Cholesky codes of Ng and Peyton [17].

The MA47 solver implements the multi-frontal sparse Gaussian elimination described in [8]. In the algorithm, the pivots used are not limited only to 1×1 diagonal pivots but also 2×2 block diagonal pivots. The solver performs a pre-factorization phase (called symbolic factorization) on the coefficient matrix to determine a pivoting order so as to minimize fill-ins. In the actual factorization process, this pivoting order may be modified to achieve better numerical stability. Note that in sparse Cholesky factorization, the pivoting order is the natural order. Because significant overhead may be incurred when the pivoting order is modified in sparse symmetric indefinite factorization, it is sometimes much more expensive than performing the sparse Cholesky factorization on a matrix with the same dimension.

The advantage of using MA47 to solve (37) is that it does not introduce a fully dense matrix in the solution process. Thus it is more suitable for SOCP problems with U having a relatively large number of columns.

Compared to the Schur complement method described above, the MA47 method does a disadvantage in that the symbolic factorization of the reduced augmented matrix need to be re-computed whenever the partition used in D changes.

7.4 Partitioning Strategy

As shown in section 6, the RAE approach for computing the search directions has the potential to overcome certain numerical instabilities encountered in the SCE approach. The RAE was derived from the augmented equation (9) by modifying the part of the coefficient matrix involving the small eigenvalues of F^2. Here we will describe the partition we use in $D = \mathrm{diag}(D_1, \bar{D}_1)$.

The choice of D_1 is dictated by the need to strike a balance between our desire to compute more accurate search directions and to minimize the size of the RAE to be solved. From the perspective of computational efficiency, it is better to have as fewer columns in the matrix U in (37) as possible, thus suggesting that the threshold for labelling an eigenvalue as "small" should be low. But from the perspective of accuracy, it is beneficial to partition eigenvalues that are smaller than, say 10^{-3}, into D_1 to improve the conditioning of the reduced augmented matrix.

With due consideration in balancing the two issues mentioned in the last paragraph, we adopt a hybrid strategy in computing the search direction in each IPM iteration.

If $\kappa(F^2) \geq 10^6$,
    put the eigenvalues of F^2 that are smaller than 10^{-3} in D_1, and the rest in $\bar{D}_1$.

Otherwise,
    put all the eigenvalues of F^2 in $\bar{D}_1$.

Some of our test problems also contain linear blocks (i.e., cones with dimensions $n_i = 1$). In this case, $F_i^2 = z_i/x_i$ is a scalar, and we put F_i^2 in D_1 if it is smaller than 10^{-3}, otherwise, we put it in $\bar{D}_1$.

As noted in Remark 3.1, when $\widetilde{A}_1$ has full row rank (for which a necessary condition is that the number of small eigenvalues put into D_1 is at least m), the Schur complement matrix M is not highly ill-conditioned, and it is not necessary to use the RAE approach to compute the search directions. When such a situation occurs, we use the SCE approach.

8 Numerical experiments

The reduced augmented equation (24) or (37) is computationally more expensive to solve than the SCE (10) because it is larger in size. As we have discussed in the last section, we can try to minimize the additional computational cost by a judicious choice of the solver used. If the number of columns in the matrix U is small, then by using the Schur complement method, the cost of solving (37) should not be much more expensive than that of solving the SCE. We adopt the following heuristic rule for the choice of the solver used for solving (37). If the number of columns in U is smaller than 200, we use the Schur complement method; otherwise, we use the MA47 method.

The RAE approach is implemented in MATLAB based on the IPM in SDPT3, version 3.1; see [26]. But the search direction in each iteration is computed based on the RAE (37).

We use the same stopping criteria mentioned in section 4. Again, the numerical results are obtained from a Pentium IV 2.4GHz PC with 1G RAM.

We consider the same SOCP problems in the section 4. But to keep our focus on the comparsion between the SCE and RAE approaches without the complication of unbounded primal solution sets, we exclude the `nbxxx`, `nqlxxx`, and `qsspxxx` problems from the numerical experiments in this section.

Our major concern in the experiments is in the following two criteria: efficiency and accuracy. We measure efficiency by the average CPU time spent per IPM iteration; while accuracy is again measured by the relative duality gap (`relgap`), and the primal and dual infeasibilities (`p-inf` and `d-inf`).

The numerical results in Table 3 show that the IPMs based on SCE may not deliver approximate optimal solutions with small primal infeasibilities. In Table 6, we see that the RAE approach can drive the primal infeasibilities of all the problems to a level of 10^{-9} or smaller. For the `schedxxx` and `randxxx` problem sets, both the SCE based IPMs in SDPT3 and SEDUMI cannot deliever accurate approximate solutions where the accuracies attained range from $\phi = -2.8$ to $\phi = -8.9$ for the `schedxxx` set and from $\phi = -4.5$ to $\phi = -8.0$ for the `randxxx` set. The IPM based on the RAE approach, however, can achieve solutions with accuracy $\phi \leq -8.8$ for all the problems in these 2 sets. The improvement in the attainable accuracy is more than 5 orders of magnitude in some cases. For the `firxxx` problems, the SCE approach can already produces accurate approximat solutions, and the RAE approach produces comparable accuracies.

The good performance in terms of accuracy of the RAE based IPM on the `schedxxx` and `randxxx` problem sets is consistent with theoretical results established in section 6. The SOCP problems in the `schedxxx` set are primal and dual nondegenerate, and strict complementarity holds at optimality. For the `randxxx` set, all the problems are primal nondegenerate, but 4 of the problems are dual degenerate. It is interesting to note that dual degeneracy does not seem to affect the performance of RAE on these degenerate problems. This fact is consistent with the observation we made in Remark 6.2.

By Theorem 5.2, the condition number of the reduced augmented matrix for the problems in `schedxxx` set is bounded when $\mu \downarrow 0$. But as noted in Remark 3.2, strict complementarity, and primal and dual nondegeneracy in an SOCP does not necessarily imply that the associated Schur complement matrix has bounded condition numbers when $\mu \downarrow 0$. The numerical results produced by the `schedxxx` set concretely show the difference in numerical stability between the SCE and RAE approaches.

From the average CPU time taken per IPM iteration for the RAE and SCE approaches in Table 6, we see that the RAE approach is reasonably efficient in that the ratio (compare with SDPT3) at most 5.0 for all the test problems, and 85% of them have ratios between 1.0 and 2.0.

The objective values obtained by the RAE based IPM are given in Table 7.

As we are able to compute rather accurate approximate solutions for (1), it is worthwhile to gather information such as primal and dual degeneracy, and strict complementarity for

some of the smaller SOCP problems we have considered in this paper. Such information is given in Table 8. We note that the degeneracies of the problems are determined by computing the numerical row and column rank (via MATLAB command `rank`) of the matrices in Theorems 20 and 21 in [1], respectively.

Table 6: A comparison between 2 SCE based IPMs and the RAE based IPM for solving SOCP problems. The last column in the table gives the maximum number of columns in the matrix U in (37).

	SDPT3		SeDuMi		RAE approach						
problem	ϕ	T_{iter}	ϕ	T_{iter}	ϕ	T_{iter}	`relgap`	`p-inf`	`d-inf`	Time	nc(U)
sched-50-50-ori	-4.3	0.19	-7.0	0.18	-10.3	0.24	4.7-11	9.4-12	1.4-11	9.3	86
sched-100-50-or	-3.8	0.42	-6.0	0.41	-10.2	0.75	6.8-11	4.8-11	1.5-11	32.4	476
sched-100-100-o	-2.8	0.82	-3.3	0.77	-8.8	1.52	1.8 -9	4.2-10	9.5-11	61.0	242
sched-200-100-o	-5.2	2.16	-3.9	1.65	-10.5	6.67	1.1-11	2.9-11	2.1-11	326.6	574
sched-50-50-sca	-6.2	0.20	-8.2	0.30	-10.6	0.30	2.7-11	1.4-12	4.3-15	9.2	85
sched-100-50-sc	-6.1	0.43	-8.9	0.46	-10.2	0.99	6.5-11	3.5-11	7.1-14	31.8	495
sched-100-100-s	-6.2	0.78	-7.1	0.66	-9.5	2.18	3.2-11	3.5-10	2.4-14	72.1	254
sched-200-100-s	-5.8	2.14	-7.8	2.39	-10.5	10.79	3.3-11	1.2-11	1.3-13	334.5	578
firL1Linfalph	-9.9	7.43	-4.7	8.83	-9.9	8.72	2.6-11	1.3-10	0.9-15	305.3	7
firL1Linfeps	-10.2	5.94	-10.4	3.32	-10.2	7.63	5.8-11	1.7-12	0.8-15	297.6	536
firL1	-10.1	27.90	-9.0	22.01	-10.1	29.37	3.3-11	7.4-11	1.0-15	675.6	0
firL2a	-10.3	4.77	-12.6	4.50	-10.3	4.99	5.0-11	6.6-16	0.8-16	39.9	804
firL2L1alph	-10.1	5.16	-3.3	4.68	-10.1	5.36	7.9-11	4.1-12	6.8-16	123.3	4
firL2L1eps	-10.4	9.92	-9.3	8.32	-10.4	14.78	3.8-11	2.3-11	0.9-15	280.9	1
firL2Linfalph	-10.1	4.25	-9.5	7.34	-10.1	7.30	7.8-11	5.6-14	7.5-16	248.2	19
firL2Linfeps	-10.2	14.62	-9.1	6.83	-10.1	15.08	7.1-11	7.9-11	6.6-16	331.7	2
firL2	-11.3	0.03	-13.1	0.03	-11.3	0.05	5.2-12	2.8-16	1.2-16	0.4	1
firLinf	-8.9	18.41	-9.3	28.50	-8.9	24.49	3.6-10	1.2 -9	1.0-15	710.3	170
rand200-300-1	-8.0	0.24	-6.4	0.76	-10.4	0.27	3.9-11	1.9-13	5.6-15	3.8	47
rand200-300-2	-6.4	0.24	-5.0	1.41	-10.3	0.29	4.5-11	6.6-14	5.6-14	4.7	62
rand200-800-1	-4.5	0.50	-5.0	2.76	-10.3	0.58	5.6-11	1.2-14	1.9-14	9.3	19
rand200-800-2	-4.1	0.52	-5.8	6.14	-10.1	0.58	7.6-11	5.1-13	7.1-14	9.9	19
rand400-800-1	-5.5	1.69	-5.1	3.23	-10.4	1.89	3.7-11	1.3-13	1.5-14	26.5	40
rand400-800-2	-5.3	1.72	-4.5	7.16	-10.6	1.90	2.5-11	3.1-13	4.7-14	28.5	40
rand700-1e3-1	-5.9	6.15	-5.7	13.45	-10.2	6.78	6.0-11	1.3-14	2.5-14	122.1	123
rand700-1e3-2	-5.5	6.22	-4.6	22.79	-10.3	6.87	4.7-11	4.7-13	6.1-14	137.4	164

Table 7: Primal and dual objective values obtained by the IPM using the RAE approach.

Problem	primal objective	dual objective
sched-50-50-ori	2.6673000977e+04	2.6673000976e+04
sched-100-50-or	1.8188993938e+05	1.8188993936e+05
sched-100-100-o	7.1736778733e+05	7.1736778607e+05
sched-200-100-o	1.4136044650e+05	1.4136044650e+05
sched-50-50-sca	7.8520384401e+00	7.8520384399e+00
sched-100-50-sc	6.7165031104e+01	6.7165031099e+01
sched-100-100-s	2.7330785593e+01	2.7330785592e+01
sched-200-100-s	5.1811961029e+01	5.1811961027e+01
firL1Linfalph	-3.0673166428e-03	-3.0673166686e-03
firL1Linfeps	-2.7112896667e-03	-2.7112897249e-03
firL1	-2.9257812784e-04	-2.9257816083e-04
firL2a	-7.1457742547e-04	-7.1457747536e-04
firL2L1alph	-5.7634915863e-05	-5.7634994782e-05
firL2L1eps	-8.4481294183e-04	-8.4481297976e-04
firL2Linfalph	-7.0591166471e-03	-7.0591167258e-03
firL2Linfeps	-1.4892049053e-03	-1.4892049762e-03
firL2	-3.1186645862e-03	-3.1186645914e-03
firLinf	-1.0068176536e-02	-1.0068176895e-02
rand200-300-1	-1.5094030119e+02	-1.5094030120e+02
rand200-300-2	-1.2861024801e+02	-1.2861024801e+02
rand200-800-1	1.8086048336e+00	1.8086048335e+00
rand200-800-2	-2.3277765218e+01	-2.3277765220e+01
rand400-800-1	6.6607764191e+00	6.6607764188e+00
rand400-800-2	6.3708631134e+01	6.3708631135e+01
rand700-1e3-1	-7.1501954797e+01	-7.1501954801e+01
rand700-1e3-2	-5.5374169004e+01	-5.5374169006e+01

Table 8: Primal and dual degeneracy, and strict complementarity information of approximate solutions of some SOCPs. A "1" means true and a "0" means false. A number of the form (34/35) in the second column means that at the computed approximate optimal solution, the column rank of $\widetilde{A}_1$ is 34, and the number of columns in $\widetilde{A}_1$ is 35.

Problem	primal non-degeneracy	dual non-degeneracy	strictly complementary
sched-50-50-orig	1	1 (79/79)	1

Table 8: Primal and dual degeneracy, and strict complementarity information of approximate solutions of some SOCPs. A "1" means true and a "0" means false. A number of the form (34/35) in the second column means that at the computed approximate optimal solution, the column rank of $\widetilde{A}_1$ is 34, and the number of columns in $\widetilde{A}_1$ is 35.

Problem	primal non-degeneracy	dual non-degeneracy	strictly complementary
sched-50-50-scaled	1	1 (83/83)	1
firL2a	1	1 (1/1)	1
firL2Linfalph	1	1 (15/15)	1
firL2	1	1 (1/1)	1
rand200_300_1	1	0 (34/35)	1
rand200_300_2	1	0 (62/65)	1
rand200_800_1	1	1 (19/19)	1
rand200_800_2	1	1 (19/19)	1
rand400_800_1	1	1 (40/40)	1
rand400_800_2	1	1 (40/40)	1
rand700_1e3_1	1	0 (84/85)	1
rand700_1e3_2	1	0 (126/130)	1

9 Conclusion

We analyzed the accuracy of the search direction computed from the SCE approach, and how the residual norm in the computed solution affects the primal infeasibility and hence the achievable accuarcy in the approximate optimal solution.

We also discussed the factors contributing to the good numerical performance of the well implemented SCE-based public domain software SEDUMI.

A reduced augmented equation is proposed to compute the search direction in each IPM iteration when the SCE cannot be solved to sufficient accuracy. The proposed reduced augmented equation approach can improve the robustness of IPM solvers for SOCP. It can be implemented efficiently by carefully preserving the sparsity structure in the problem data. Numerical results show that the new approach can produce approximate optimal solutions that are much more accurate than that produced by the SCE approach.

Acknowledgement

The authors are grateful to Professor Robert Freund for valuable suggestions on the manuscript.

References

[1] F. Alizadeh and D. Goldfarb, *Second-order cone programming*, Mathematical Programming, 95 (2003), pp. 3–51.

[2] F. Alizadeh and S.H. Schmieta, *Optimization with Semi-definite, Quadratic and Linear Constraints*, Report 23-97, Rutgers Center for Operations Research, Rutgers University, 1997. Available from `http://rutcor.rutgers.edu/pub/rrr/reports97/23.ps`

[3] E.D. Andersen, C. Roos, and T. Terlaky, *On implementing a primal-dual interior-point method for conic quadratic optimization*, Mathematical Programming, Series B, 95 (2003),pp. 249-277.

[4] K.D. Andersen, *A modified Schur complement method for handling dense columns in interior point methods for linear programming*, ACM Transactions on Mathematical Software, 22 (1996), pp. 348–356.

[5] S.C Chan, K.M. Tsui, and K.S. Yeung, *Design of FIR digital filters with prescribed flatness and peak error constraints using second order cone programming*, preprint, Department of Electrical and Electronic Engineering, University of Hong Kong, Hong Kong.

[6] H. Ciria and J. Peraire, *Computation of upper and lower bounds in limit analysis using second-order cone programming and mesh adaptivity*, 9th ASCE Specialty Conference on Probabisitic Mechanics and Structural Reliability, 2004.

[7] T.A. Davis, *UMFPACK version 4.0 user guide*, Department of Computer and Information Science and Engineering, U. of Florida, Gainesville, Florida.

[8] I.S. Duff and J.K. Reid, *The multifrontal solution of indefinite sparse symmetric linear equations*, ACM Transactions on Mathematical Software, 9 (1983), pp. 302–325.

[9] A. George, and K.H. Ikramov, *On the condition of symmetric quasi-definite matrices*, SIAM J. Matrix Analysis and Applications, 21 (2000), pp. 970–977.

[10] D. Goldfarb, and K. Scheinberg, *A product-form Cholesky factorization implementation of an interior-point method for second order cone programming*, preprint, IEOR, Columbia University, NY.

[11] G. Golub and C. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, USA, 1996.

[12] M. Gu, *Primal-dual interior-point methods for semidefinite programming in finite precision*, SIAM J. Optimization, 10 (2000), pp. 462–502.

[13] N.J. Higham, *Accuracy and stability of numerical algorithms*, SIAM, Philadelphia, 1996.

[14] V. Kovacevic-Vujcic and M.D. Asic, *Stablization of interior-point methods for linear programming*, Computational Optimzation and Applications, 14 (1999), pp. 331–346.

[15] M. S. Lobo, L. Vandenberghe, S. Boyd and H. Lebret, *Applications of Second-order Cone Programming*, Linear Algebra Appl., 284 (1998), pp.193–228.

[16] R.D.C. Monteiro and T. Tsuchiya, *Polynomial Convergence of Primal-Dual Algorithms for the Second-Order Cone Program Based on the MZ-Family of Directions*, Mathematical Programming, 88 (2000), pp. 61–83.

[17] E.G. Ng and B.W. Peyton, *Block sparse Cholesky algorithms on advanced uniprocessor computers*, SIAM J. Scientific and Statistical Computing, 14 (1993), pp. 1034–1056.

[18] G. Pataki and S.H. Schmieta, *The DIMACS Library of Mixed Semidefinite-Quadratic-Linear Programs*,
`http://dimacs.rutgers.edu/Challenges/Seventh/Instances/`

[19] J. Reid, and I.S. Duff, *MA47, a Fortran code for direct solution of indefinite sparse symmetric linear systems*, Report RAL-95-001, Rutherford Appleton Laboratory, Oxfordshire, England, January, 1995.

[20] Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Company, Boston, 1996.

[21] D. Scholnik and J. Coleman, *An FIR Filter Optimization Toolbox for Matlab 5 & 6*, available from `http://www.csee.umbc.edu/~dschol2/opt.html`

[22] 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.

[23] J.F. Sturm, *Implementation of interior point methods for mixed semidefinite and second order cone optimization problems*, Optimization Methods and Software, 17 (2002), pp. 1105–1154.

[24] J.F. Sturm, *Avoiding numerical cancellation in the interior point method for solving semidefinite programs*, Mathematical Programming, Series B, 95 (2003), pp. 219–247.

[25] K.C. Toh, M.J. Todd, and R.H. Tütüncü, *SDPT3- A Matlab Software package for Semidefinite Programming*, Optimization Methods and Software, 11 (1999), pp. 545–581.

[26] R.H. Tütüncü, K.C. Toh and M.J. Todd, *Solving Semidefinite-Quadratic-Linear Programming Using SDPT3*, Mathematical Programming, Series B, 95 (2003), pp. 189–217.

[27] K.C. Toh, R.H. Tütüncü, and M.J. Todd, *On the implementation of SDPT3 (version 3.1) – a* MATLAB *software package for semidefinite-quadratic-linear programming*, invited paper, 2004 IEEE Conference on Computer-Aided Control System Design, Taipei, Taiwan.

[28] R.J. Vanderbei, *Symmetric Quasidefinite Matrices*, SIAM J. Optimization, 5 (1995), pp. 100–113.

[29] S.J. Wright, *Stability of linear equations solvers in interior-point methods*, SIAM J. Matrix Analysis and Applications, 16 (1995), pp. 1287–1307.

[30] S.J. Wright, *Stability of augmented system factorizations in interior-point methods*, SIAM J. Matrix Analysis and Applications, 18 (1997), pp. 191–222.

[31] Y. Ye, M.J. Todd, and S. Mizuno, *An $O(\sqrt{n}L)$-iteration homogeneous and self-dual linear programming algorithm*, Math. Oper. Res., 19 (1994), pp. 53–67