The Gauss-Newton Direction in Semidefinite Programming

Serge Kruk * Masakazu Muramatsu † Franz Rendl ‡
Robert J. Vanderbei § Henry Wolkowicz ¶

May 22, 1998
University of Waterloo
Department of Combinatorics and Optimization
Waterloo, Ontario N2L 3G1, Canada
Research Report CORR 98-16

Key words: Semidefinite Programming, Newton Direction, Stability, Symmetrization.

Abstract

Primal-dual interior-point methods have proven to be very successful for both linear programming (LP) and, more recently, for semidefinite programming (SDP) problems. Many of the techniques that have been so successful for LP have been extended to SDP. In fact, interior point methods are currently the only successful techniques for SDP.

---

*Research supported by Natural Sciences Engineering Research Council Canada. E-mail sgkruk@acm.org

†Department of Mechanical Engineering, Sophia University, 7-1 Kioi-cho, Chiyoda-ku, Tokyo 102 Japan

‡Technische Universität Graz, Institut für Mathematik, Steyrergasse 30, A-8010 Graz, Austria

§EMS Program Director, ACE-42 E-Quad, Princeton University, Princeton NJ 08544, Tel: 609-258-0876, Fax: 609-258-3796, E-mail rvdb@princeton.edu, http://www.princeton.edu/~rvdb/

¶Research supported by Natural Sciences Engineering Research Council Canada. E-mail hwolkowi@orion.math.uwaterloo.ca, http://orion.math.uwaterloo.ca/~hwolkowi

[0]This report is available by anonymous ftp at orion.uwaterloo.ca, in directory pub/henry/reports and also with URL http://orion.uwaterloo.ca/~hwolkowi/henry/reports/gnsdp.ps.gz

We present a new paradigm for deriving these methods: 1) using the optimality conditions from the dual log-barrier problem, we obtain primal feasibility, dual feasibility, and perturbed complementary slackness equations; 2) the perturbed complementary slackness condition is quite nonlinear, so we modify this condition to obtain a bilinear condition, i.e. a condition that is less nonlinear; 3) we now find a search direction by applying the Gauss-Newton method to the least squares problem for these optimality conditions; equivalently we find the least squares solution of the linearized perturbed optimality conditions.

In the case of LP, the Gauss-Newton direction for the least squares problem is equivalent to the Newton direction applied to solving the modified (square) nonlinear system of optimality conditions. Though this paradigm does not directly provide a new search direction for linear programming, it does provide a new approach for convergence proofs as well as motivation for step lengths larger than one.

However, there is one major difference between LP and SDP that raises several interesting questions. That difference is the form of the perturbed complementarity condition used in the optimality conditions. In LP this condition is modified to be $ZX - \mu I = 0$. The primal matrix X and the dual slack matrix Z are diagonal in LP but may only be symmetric in SDP; this results in ZX not being symmetric in general, i.e. the optimality conditions in the SDP case end up as an overdetermined system of nonlinear equations.

There have been various approaches which modify the complementarity condition so that the linearization of the optimality conditions are "square", i.e. map between the same spaces. These approaches can have several drawbacks, e.g. numerical instability near the optimum and lack of positive definiteness after symmetrization. Our least squares approach requires no symmetrization and does not suffer from these drawbacks. We concentrate on solving the overdetermined, system in the best way possible. In particular, we use Gauss-Newton type methods. This leads to numerically stable as well as excellent search directions which lead to the central path. Though the numerical efficient calculation of the Gauss-Newton direction is still an open question, we present a preliminary "top down" elimination approach that is competitive timewise and empirical evidence suggests that it is often more robust than other directions currently in use.

Contents

1 INTRODUCTION

Primal-dual interior-point (p-d i-p) methods have proven to be very successful for both linear programming (LP) and, more recently, for semidefinite programming (SDP) problems. Many of the techniques that have been so successful for linear programming have been extended to SDP. In fact, interior-point methods currently appear to be the only successful techniques for SDP.

The derivation of interior-point algorithms for SDP has followed closely upon the lessons learned from the very successful applications in LP. This has led to specific choices for the form of the perturbed optimality conditions to which Newton's method is applied. In this paper, we present a new paradigm for deriving these methods both for LP and for SDP.

1. Using the optimality conditions from the dual log-barrier problem, we obtain the system of three equations: primal feasibility; dual feasibility; and perturbed complementary slackness, see Section 1.2, program (1.2) and equations (1.4), (1.5), and (1.6).

2. The linear system for the Newton direction (N-direction) for the log-barrier problem is ill-conditioned but can be solved in a stable way by using an augmented system, i.e. this direction itself is *not ill-conditioned* and can be used successfully. However, the perturbed complementary slackness condition (1.6) is quite nonlinear, so we modify this condition to obtain a bilinear condition, i.e. a condition that is less nonlinear, since Newton type methods converge much faster on systems that are less nonlinear.

3. We now find the Gauss-Newton search direction (GN-direction) to the least squares problem for the resulting equations; equivalently we find the least squares solution of the linearized perturbed optimality conditions.

This paradigm is very similar to what is currently done, e.g. see [12]. In fact, item 1 is unchanged. However, the motivation for the modifications of the optimality conditions

in item 2 is usually from the ill-conditioning. Here we emphasize that this ill-conditioning can be avoided (see [26] Section 5) and that the true motivation should be to obtain a less nonlinear system. The main difference with our paradigm is item 3, where we use a GN direction on the least squares problem rather than a Newton direction on modified optimality conditions. This approach shows that the search direction for SDP can be found from the optimality conditions just as is done in LP. This is in contrast to current approaches which modify the optimality conditions in order to be able to apply a Newton method to a square system. Our GN approach gives rise to two new stable and efficient directions.

The p-d i-p approach has proven to be the superior choice of interior-point methods for LP. The current favourite is the Mehrotra predictor-corrector approach. This can be viewed as a simplified, composite, damped Newton method on the optimality conditions, with a backtracking linesearch, see [12]. In addition, it is shown in [12] that the method cannot be viewed as a Newton method on the log-barrier problem but can be explained as a modification of the optimality conditions that removes the ill-conditioning that arises in log-barrier Newton methods. However, as discussed in the above paradigm, the ill-conditioning is a *red herring*. The modification reduces the nonlinearity of the complementarity condition and thus speeds up Newton type methods.

1.1 Outline

In this paper we introduce and study the properties of search directions for primal-dual interior-point methods for semidefinite programming, where these directions arise from a Gauss-Newton method applied to minimizing the norm of the (overdetermined) system of optimality conditions, rather than a Newton method applied directly to modified optimality conditions. In addition, we compare several directions to see which direction results in fewer iterations in the p-d i-p framework. We conclude that the Gauss-Newton method is competitive and often more robust.

We complete this section with the background on SDP, see Section 1.2. Then in Section 2 we define the GN direction. Section 2.1.1 contains some theoretical results on the GN direction, e.g. results on full rank of the linearized systems.

The solution techniques are discussed in Section 3. Section 3.1 contains the details of the elimination approach used in our implementation of the GN direction. In addition, we show that the linear system for the search direction is full rank at each iteration as well as in the limit. (This is not true for the various directions in use at this time.)

The numerical tests appear in Section 4. These tests compare various search directions. We see that our new direction is very competitive in number of iterations, cpu time, and robustness. In particular, it is superior when we work with certain badly scaled problems.

Remark 1.1 *In the detailed web version of the present paper [26, Section 5], we recall how ill-conditioning arises from log-barrier problems in NLP and how this can be avoided using an augmented system. In the last part of [26] we deal with justifying the ill-conditioning claims for our paradigm.*

We conclude with several remarks and open questions in Section 5.

1.2 Background

We now include the notation and background information on SDP. Comprehensive upto-date notation is available on WWW with URL:
http://orion.uwaterloo.ca/~hwolkowi/henry/software/psd_tool.d/sdpnotation.d/notation.ps.

Let $\mathcal{S}_n$ denote the space of $n \times n$ symmetric matrices equipped with the trace inner product, $\langle A, B \rangle = \operatorname{trace} AB$; and let $A \succeq 0$ (resp. $A \succ 0$) denote positive semidefiniteness (resp. positive definiteness), $A \succeq B$ denotes $A - B \succeq 0$, i.e. $\mathcal{S}_n$ is equipped with the Löwner partial order. We let $\mathcal{P}$ denote the cone of symmetric positive semidefinite matrices; $\mathcal{M}_n$ denotes the set of $n \times n$ real matrices.

SDP is an extension of linear programming, denoted LP, where nonnegative variables are replaced by semidefinite matrices. Our (primal) semidefinite programming problem is

$$(PSDP) \qquad \begin{array}{rlll} \mu^* := & \max & \operatorname{trace} CX \\ & \text{subject to} & \mathcal{A}X = b \\ & & X \succeq 0, \end{array}$$

where the linear operator $\mathcal{A} : \mathcal{S}_n \to \Re^m$ is defined as

$$\mathcal{A}X = (\operatorname{trace} A_i X),$$

and $A_i,\ i = 1, \ldots m$, are given symmetric matrices. To avoid trivialities, we assume that μ^* is finite valued. Thus the problem has a bounded optimal value and a feasible solution. We also assume that the matrices $\{A_i : i = 1, \ldots, m\}$ are linearly independent.

This problem has recently generated a lot of interest. One reason is that there are many diverse applications: in discrete optimization see e.g. [3, 18, 41]; in engineering see e.g. [9, 56]; for matrix completions see e.g. [21, 1]. Another reason for the interest is that SDPs can be solved efficiently using interior-point methods. More applications and evidence of the current high level of activity can be found in the recent theses: [2, 40, 19, 36, 22, 60, 25, 23, 50]; and in the recent books and notes [58, 33, 57, 52, 44, 35].

We now follow the p-d i-p approach developed in [20]. We follow the usual derivation of p-d i-p methods in Linear Programming and first introduce the associated *log-barrier*

problem for the dual SDP

$$(DSDP) \qquad \begin{array}{ll} \min & \operatorname{trace} b^t y \\ \text{subject to} & \mathcal{A}^* y - Z = C \\ & Z \succeq 0, \end{array} \tag{1.1}$$

where $\mathcal{A}^*$ denotes the *adjoint operator*, $\mathcal{A}^* y = \sum y_i A_i$. The *dual barrier problem* is:

$$(DBP) \qquad \begin{array}{ll} \text{minimize} & b^t y - \mu \log\det Z \\ \text{subject to} & \mathcal{A}^* y - Z = C \\ & Z \succeq 0, \end{array} \tag{1.2}$$

where μ is a positive real number called the *barrier parameter*.

For each $\mu > 0$, there is a corresponding Lagrangian:

$$L_\mu(X, y, Z) = \quad b^t y - \mu \log\det Z + \operatorname{trace} X\,(Z + C - \mathcal{A}^* y)\,. \tag{1.3}$$

With $X_\mu, Z_\mu \succ 0$, the first-order optimality conditions for the saddle point of this Lagrangian are obtained by differentiation:

$$\nabla_X L_\mu = Z + C - \mathcal{A}^* y = 0 \tag{1.4}$$

$$\nabla_y L_\mu = b - \mathcal{A}(X) = 0 \tag{1.5}$$

$$\nabla_Z L_\mu = X - \mu Z^{-1} = 0. \tag{1.6}$$

There exists a unique solution (X_μ, y_μ, Z_μ) to these optimality conditions provided both the primal and dual SDPs satisfy the Slater constraint qualification, i.e. strict feasibility. The one-parameter family $\{(X_\mu, y_\mu, Z_\mu) : 0 \le \mu \le \infty\}$ is called the *central trajectory* or the central path. Given a point (X, y, Z) on the central trajectory it is easy to determine its associated μ value using (1.6):

$$\mu = \frac{\operatorname{trace}(ZX)}{n}. \tag{1.7}$$

We note that if the point is a feasible solution of the primal and the dual problem, then $\operatorname{trace}(ZX)$ is the gap between the primal and the dual objective value.

For simplicity of notation, we rewrite the optimality conditions as the function

$$F_\mu(s) = F_\mu(X, y, Z) := \begin{pmatrix} Z + C - \mathcal{A}^* y \\ b - \mathcal{A}(X) \\ X - \mu Z^{-1} \end{pmatrix} =: \begin{pmatrix} F_d \\ F_p \\ F_{Z,X} \end{pmatrix}. \tag{1.8}$$

The equations correspond to: dual feasibility, primal feasibility, and complementary slackness, respectively. The solution s^* to $F_\mu(s) = 0$ satisfies the Karush-Kuhn-Tucker conditions (1.4) to (1.6) and is the optimal solution to the barrier problem.

SDP has many similarities with LP. One important point, mentioned above, is that SDP problems can be solved very efficiently using p-d i-p methods. These methods are based on applying Newton's method to the Karush-Kuhn-Tucker (KKT) optimality conditions. Many of the algorithms from LP can be extended to SDP and retain many of their nice properties. (See e.g. the books on interior-point methods: by Nesterov and Nemirovsky [33]; by Wright [57]; by Ye [58]; and edited by Terlaky [52].)

However, there are several subtle differences with LP.

- One interesting difference is that duality gaps can exist for SDP in the absence of strictly feasible solutions. This means that strong duality can fail, i.e. the primal or dual may not be attained and/or there is a nonzero duality gap between the primal and dual optimal values. However, feasible interior-point methods usually assume strictly feasible solutions, so we will not dwell on that here. (See e.g. [39] for a discussion on strong duality and [37] for the generic properties.)

- Another difference is that strict complementarity may fail at the optimum, i.e. there may be no optimal primal-dual pair for which $X + Z$ is full rank. (Here Z is the dual slack variable.) (See e.g. [5, 45] for a discussion on strict complementarity and [37] for the generic properties.)

$\rightarrow \star$ Finally, the difference that we focus on is that the perturbed complementary slackness equations, that are essential for p-d i-p methods, involve possibly nondiagonal matrices and can be written in many different ways.

For example, if Z is the slack matrix in the dual (1.1) then:

1. the following comes from the optimality conditions for the log-barrier problem of the dual SDP problem (1.2)
$$X - \mu Z^{-1} = 0;$$

2. a nonlinear transformation reduces the nonlinearity and yields
$$ZX - \mu I = 0;$$

3. equivalent forms for the above two equations are
$$Z - \mu X^{-1} = 0,$$
$$XZ - \mu I = 0;$$

4. a general symmetrization operator yields, see e.g. [30],

$$H_P(M) := \frac{1}{2}\left(PMP^{-1} + (PMP^{-1})^t\right) \tag{1.9}$$

where $M = XZ$ and P is any nonsingular matrix;

   (a) the A.H.O. direction [6] corresponds to choosing $P = I$; this is equivalent to symmetrizing ZX, i.e. replacing it by $(ZX + XZ)/2$;

   (b) the H.K.M. direction, ([20], [24], [27]), corresponds to choosing $P = Z^{1/2}$; this is equivalent to symmetrizing ΔX after solving the linear system with ΔX treated as a general matrix;

   (c) the NT direction, [55], corresponds to $P = G^{-1}$, where G is the unique matrix such that $D := G^t Z G = G^{-1} X G^{-t}$ is a diagonal matrix.

(See [6, 53, 44] for comparisons of various Newton directions.)

For SDP, there is an additional complication that arises from the various forms of the complementarity conditions. The most popular form of these conditions for LP is $ZX - \mu I = 0$. In the case of SDP, as mentioned above, this form does not map between the same spaces since ZX is not symmetric unless Z, X commute. (It is symmetric on the central path, i.e. when $ZX = \mu I$.) Moreover, the derivative (linearization) of the optimality conditions may not be a one-one map since strict complementarity may fail. The above equivalences present several of the approaches that "symmetrize" the complementarity conditions in order to get a "square" system so that they can apply Newton's method for the search direction, see e.g. [42, 20, 27, 24, 30, 29, 6, 59, 55, 47]. In fact, to our knowledge all of the current SDP p-d i-p algorithms perform some sort of symmetrization. (See e.g. the software packages [54, 16, 10, 4].)

2 THE ZX FORM of COMPLEMENTARITY

We now study the most popular form of the complementarity conditions i.e.

$$ZX - \mu I = 0. \tag{2.1}$$

This is the one that is used in LP. Though many other linearizations appear to be based on different forms, e.g. the symmetrization $XZ + ZX - 2\mu I = 0$ studied in [5, 6], these are all still trying to solve the above form. Thus, the main question is how to solve this form efficiently.

We now look at (2.1) as an operator mapping the symmetric matrix space $\mathcal{S}_n \times \mathcal{S}_n$ to the space of matrices $\mathcal{M}_n$. With this in mind, we see that we have an overdetermined

system of nonlinear equations. The standard approach, e.g. [11], to solving an overdetermined nonlinear system is to solve an equivalent least squares problem, see (2.3), i.e. no symmetrization is needed. We now explore this possibility and see that there are many advantages for this point of view.

Equivalently, we are solving the equation

$$\Phi_Z F_\mu(s) = \Phi_Z F_\mu(X,y,Z) = \begin{pmatrix} Z + C - \mathcal{A}^* y \\ b - \mathcal{A}(X) \\ ZX - \mu I \end{pmatrix} =: \begin{pmatrix} F_d \\ F_p \\ F_c \end{pmatrix} = 0, \tag{2.2}$$

where F_μ is defined in (1.8), Φ_Z is the block diagonal operator Diag $(I\ I\ Z)$, and $\cdot^*$ denotes adjoint operator. Thus we are solving a modified, less nonlinear, system. (Such modifications can result in improved speed of convergence, see e.g. [7, 17, 38]. In fact, such a modification is the basis behind the success of algorithms for trust region subproblems, see e.g. [38, 49, 43, 31].)

2.1 The GN-Direction

The standard approach to solving a system of nonlinear equations is to apply the Gauss-Newton method, i.e. we minimize the sum of squares of the equations and use an approximation to the Newton search direction, the GN direction.

Under very mild assumptions, e.g. full column rank of the derivative $(\Phi_Z F_\mu)'$, (which holds here, Theorem 2.2) there are several advantages to using a Gauss-Newton method for solving nonlinear equations, see e.g. [11, 13, 15, 14].

1. If the optimal value of the least squares problem is 0, then we get quadratic convergence asymptotically.
   In our case, if we assume that the duality gap is 0 and that both primal and dual SDP are attained, then there is a solution for the optimality conditions which means that the least squares problem does indeed have a (attained) minimum value of 0 for each $\mu > 0$.

2. The GN direction always exists.
   Thus there are no problems with existence as there are for the other search directions used in the literature, see e.g. [48, 55, 29] for the AHO direction and [28] for the MT (Monteiro-Tsuchiya) direction. (In fact, one could apply the Gauss-Newton approach to the optimality conditions used for other directions and thus avoid the question of existence.)

3. The direction is always a descent direction for the least squares problem.
   Thus the least squares problem can be used as a merit function for monitoring sufficient decrease. (Though it is not a true potential function, see [32, page 5].)

4. The direction is scale invariant.

5. The search direction is found using the same equations as for LP, with the exception that a least squares solution is used.

With this motivation, we now look at the nonlinear least squares problem

$$R_\mu^* = \min_s f_\mu(s) := \frac{1}{2}||\Phi_Z F_\mu(s)||_2^2. \tag{2.3}$$

The optimal value R_μ^* is called the residual.

Remark 2.1 *The GN direction is formed by dropping the second order terms of the Hessian of f_μ from the Newton equation $\nabla^2 f_\mu(s) d_N = -\nabla f_\mu(s)$. Recall that $F_c = F_c(X,Z) = ZX - \mu I$. The second order terms are: $\langle F_c(Z,X), F_c''(Z,X)(\cdot,\cdot)\rangle$. If we denote the changes in Z, X as l, h respectively, then*

$$(Z+l)(X+h) - \mu I = ZX - \mu I + Zh + lX + lh.$$

The curvature in this direction is therefore:

$$\text{trace}\,(ZX - \mu I)lh.$$

In fact, the second order terms that are dropped from the Newton direction to get the GN direction are rather simple, i.e. similar to the LP case

$$H := \frac{1}{2}\begin{bmatrix} 0 & (ZX-\mu I) \\ (ZX-\mu I) & 0 \end{bmatrix}.$$

This matrix is indefinite and can cause problems if included in the approximate Hessian. Note that the residual $R_\mu^* = 0$ in our case, which implies that the matrix of second derivatives is 0 at the optimum in (2.3). Thus the error is very small near the central path.

Finding the GN-direction is therefore equivalent to solving the normal equations

$$\left((\Phi_Z F_\mu)'\right)^* (\Phi_Z F_\mu)'(\Delta s) = -\left((\Phi_Z F_\mu)'\right)^* (\Phi_Z F_\mu), \tag{2.4}$$

which is equivalent to finding the least squares solution of the overdetermined system

$$(\Phi_Z F_\mu)'(\Delta s) = -(\Phi_Z F_\mu). \tag{2.5}$$

One popular method for solving the above system is to solve (2.5) using a (sparse, e.g. [8, 51]) QR decomposition. The N-direction for the least squares problem is solved from the system

$$\left(\left((\Phi_Z F_\mu)'\right)^* (\Phi_Z F_\mu)' + H\right)(\Delta s) = -\left((\Phi_Z F_\mu)'\right)^* (\Phi_Z F_\mu) \tag{2.6}$$

which cannot exploit the normal equations. Moreover, the addition of the operator H introduces negative eigenvalues and so the N-direction may not point to the minimum.

The overdetermined linear system we are solving is:

$$\begin{bmatrix} 0 & -\mathcal{A}^* & I \\ -\mathcal{A} & 0 & 0 \\ Z\cdot & 0 & \cdot X \end{bmatrix} \begin{pmatrix} \Delta X \\ \Delta y \\ \Delta Z \end{pmatrix} = \begin{pmatrix} -F_d \\ -F_p \\ -F_c \end{pmatrix}. \tag{2.7}$$

This system yields the GN-direction. The system is very sparse and most of the terms do not change at each iteration. We can exploit this fact as is done in linear programming where block elimination is used.

2.1.1 Theoretical Properties

We have assumed that the operator $\mathcal{A}$ is onto, i.e. the matrices A_k are linearly independent or, equivalently, $\mathcal{A}^*$ is one to one or full column rank. We now see that we get existence of the GN-direction without any other assumptions. However, to guarantee full rank of the Jacobian at the optimum, we do need to assume that the primal and dual optimum are both unique. This is equivalent to primal and dual nondegeneracy as defined and used in [5, 6] and shown to be generic, see also [46, 37].

We first prove the well known property of uniqueness of points on the central path.

Theorem 2.1 *Suppose that (X, y, Z) lie on the central path with log-barrier parameter $\mu > 0$. Then 0 is the unique solution of the system (2.7).*

Proof. Suppose that $(\Delta X, \Delta y, \Delta Z)$ solves (2.7). The central path properties yield

$$\Delta Z = \mathcal{A}^*(\Delta y),\ \mathcal{A}\Delta X = 0,\ Z\Delta X + \mathcal{A}^*(\Delta y)X = 0,\ Z^{-1} = \frac{1}{\mu}X.$$

Therefore

$$\begin{aligned} 0 &= \mathcal{A}\Delta X \\ &= -\mathcal{A}(Z^{-1}\mathcal{A}^*(\Delta y)X) \\ &= -\frac{1}{\mu}\mathcal{A}(X\mathcal{A}^*(\Delta y)X) \\ &= -\left(\frac{1}{\mu}\sum_i (\Delta y)_i \,\text{trace}\,((A_kX)(A_iX))\right) \\ &= -\left(\frac{1}{\mu}\sum_i (\Delta y)_i \,\text{trace}\,\left((X^{\frac{1}{2}}A_kX^{\frac{1}{2}})(X^{\frac{1}{2}}A_iX^{\frac{1}{2}})\right)\right). \end{aligned}$$

The $m \times m$ matrix

$$\left[\text{trace}\,(X^{\frac{1}{2}}A_kX^{\frac{1}{2}})(X^{\frac{1}{2}}A_iX^{\frac{1}{2}})\right]$$

is a Gram matrix and is nonsingular since the matrices A_k, and so also $X^{\frac{1}{2}}A_kX^{\frac{1}{2}}$, are linearly independent. This implies that $\Delta y = 0$.

□

We can also prove that the linear system has full rank off of the central path, which implies that the GN-direction always exists.

Theorem 2.2 *Suppose that both X and Z are positive definite. Then the linear operator in system (2.7) has full rank.*

Proof. Suppose that $(\Delta X, \Delta y, \Delta Z)$ solves (2.7) with right hand side 0. Then $\Delta Z = \mathcal{A}^*\Delta y$. Therefore, $\mathcal{A}\Delta X = 0$ and

$$Z\Delta XX^{-1} = S := -\mathcal{A}^*\Delta y \perp \Delta X. \tag{2.8}$$

This means that the quadratic form

$$0 = \text{trace}\, S\Delta X = \text{trace}\, Z\Delta XX^{-1}\Delta X = ||Z^{\frac{1}{2}}\Delta XX^{-\frac{1}{2}}||_F^2, \tag{2.9}$$

i.e. necessarily $\Delta X = 0$. This means that both Δy and ΔZ are 0 as well.

□

Under strict complementarity and nondegeneracy, we can also prove that the linear system at optimality is full rank. This, in combination with Theorem 2.2 above, implies that we do not run into ill-conditioning difficulties.

Theorem 2.3 *Suppose that (X, y, Z) are a unique optimal primal and dual solution (equivalently primal and dual nondegenerate) of PSDP. In addition, suppose that strict complementarity holds, i.e. $X + Z \succ 0$. Then 0 is the unique solution of the system (2.7) at $\mu = 0$.*

Proof. Suppose that $(\Delta X, \Delta y, \Delta Z)$ solves (2.7). Then $\Delta X, \Delta Z$ are symmetric and so $(\Delta X, \Delta y, \Delta Z)$ also solves the symmetrized system

$$Z\Delta X + \Delta ZX + \Delta XZ + X\Delta Z = 0.$$

The result in [6, Theorem 1] implies that the Jacobian corresponding to the symmetrized system is nonsingular and so necessarily $(\Delta X, \Delta y, \Delta Z) = 0$.

We can provide a self-contained proof using the approach above in Theorem 2.2, i.e. if $w = \Delta X$ and $s = \mathcal{A}^*\Delta y$, then $Zw + sX = 0 = ZX$, and $Z + X \succ 0$. Since Z and X commute, we can assume that they are diagonal. The strict complementary slackness and

uniqueness implies that both s and w are 0.

□

Another important property is invariance.

Theorem 2.4 *The GN direction, the least-square solution to (2.5) used to solve program (2.3), is invariant under affine transformation of the variables $r = Hs + h$ where H is non-singular.*

Proof. The GN step in s-space, from current point s_c, is given by

$$s_+ = s_c - \left((\Phi_Z F_\mu(s_c))'\right)^\dagger (\Phi_Z F_\mu(s_c)) .$$

Now, under the affine scaling $r = Hs+h$ we obtain that $\Phi_Z F_\mu(s) = \Phi_Z F_\mu(H^{-1}(r-h)) =: G(r)$. Using the full rank property in Theorem 2.3, we see that the GN step in the r-space, from the current point $r_c = Hs_c + h$, will be

$$\begin{aligned} r_+ &= r_c - (G(r_c)')^\dagger G(r_c) && (2.10)\\ &= Hs_c + h - \left((\Phi_Z F_\mu(s_c))' H^{-1}\right)^\dagger (\Phi_Z F_\mu(s_c)) && (2.11)\\ &= Hs_c + h - H\left((\Phi_Z F_\mu(s_c))'\right)^\dagger (\Phi_Z F_\mu(s_c)) && (2.12)\\ &= Hs_+ + h. && (2.13) \end{aligned}$$

The GN step is therefore invariant under such transformations of the space.

□

Perhaps these properties explain the good behavior of the GN step for the badly scaled problems described in Section 4.

Remark 2.2 *We can find redundant equations in (2.2) and solve a smaller system. We can remove the lower triangular part from the nonlinear equation, i.e. we get the equivalent equation*

$$\mathcal{T}(ZX - \mu I) = 0, \qquad (2.14)$$

where the linear operator $\mathcal{T} : \mathcal{M}_n \to \mathcal{S}_n$, by ignoring the strictly lower triangular part of the matrix, i.e. the i, j components are

$$(\mathcal{T}(W))_{ij} = \begin{cases} W_{ij} & \text{if } i \le j \\ W_{ji} & \text{if } i > j \end{cases} .$$

(Note $\mathcal{T}(M) = PMP^{-1}$ implies $P^{-1}\mathcal{T}(M)P = M$ which shows that this does not fall under the general symmetrization framework (1.9).)

The resulting system is now square, i.e. the domain and range space are the same. Note that $\mathcal{T}$ is the composition of two linear maps, i.e. first apply the simple coordinate (orthogonal) linear projection onto the subspace of upper triangular matrices in the space of $n \times n$ matrices; then apply the one-one mapping to the space of symmetric matrices formed by duplicating the upper triangular part in the lower triangular part. We have $\mathcal{T}(\mathcal{T}(W)) = \mathcal{T}(W)$, i.e. $\mathcal{T}$ is idempotent (a projection). In fact, if we identify the upper triangular matrices with $\mathcal{S}_n$, then we can consider $\mathcal{T}$ as orthogonal as well.

The optimality conditions become

$$\begin{pmatrix} Z + C - \mathcal{A}^* y \\ b - \mathcal{A}(X) \\ \mathcal{T}(ZX - \mu I) \end{pmatrix} = 0. \tag{2.15}$$

This does not introduce new nonlinearities but does guarantee that we map between the same spaces. In fact, we can prove that we do not lose information in the optimality conditions when we only consider the upper triangular parts.

Lemma 2.1 *Suppose that $Z \succ 0$. Then the linear operator $\mathcal{T}_Z(\cdot) := \mathcal{T}(Z\cdot)$ is a one-one mapping on $\mathcal{S}_n$. In addition, if $\mathcal{T}_Z(X) = \mu I$, for some $X = X^t$, then $ZX = \mu I$.*

Though we do not include numerical tests with this direction, the tests that we have done show that this direction is very robust and does not suffer from ill-conditioning near the optimum. (For more details see [26].)

3 LINEARIZED LEAST SQUARES

The efficient calculation of our GN direction is still an open problem. This has to involve a clever implementation of a QR algorithm. Therefore, we present a constrained GN, denoted RGN, direction that is easier to calculate, see (3.3). The RGN direction still maintains the nice properties of the GN direction.

3.1 Basic Elimination

In order to reduce the execution time of the least-square solver, we can algebraically eliminate some variables using block elimination. The order we eliminate is different than that used in LP or other SDP codes, where ΔX is eliminated first from the linearized complementarity condition. One advantage that we have is that we never invert Z and so do not appear to get numerical problems when we are very close to the optimum.

It is simplest to first force dual feasibility and eliminate ΔZ from (2.7), which, we recall, is of size $(t(n) + m + n^2) \times (t(n) + m + t(n))$. We use

$$\Delta Z = \mathcal{A}^*(\Delta y) - F_d,$$

to obtain the $(m + n^2) \times (t(n) + m)$ reduced *dual feasible* system

$$\begin{bmatrix} -\mathcal{A} & 0 \\ Z\cdot & \mathcal{A}^* \cdot X \end{bmatrix} \begin{bmatrix} \Delta X \\ \Delta y \end{bmatrix} = \begin{bmatrix} -F_p \\ F_d X - F_c \end{bmatrix}. \tag{3.1}$$

The upper part of this equation is still sparse and can therefore be treated separately once for all iterations. Moreover, since the backsolve for ΔZ is a sum of matrices it is inexpensive and exploits any sparsity in $\mathcal{A}^*$.

A smaller system, everything else being equal, is faster to solve than a large one. In the absence of specific information about the structure of the problem, we pursue this idea further since it is possible to eliminate more variables and obtain an even smaller system. For a specific problem with known sparsity structure, maxcut for example, it is possible to eliminate variables while maintaining sparsity. In the general case we discuss here, the major bottleneck being the size of the system, we aim at a reduction in size with little concern for sparsity.

We now consider the primal constraint operator $\mathcal{A}$, which we write as

$$\mathcal{A} = [\mathcal{A}_B \mathcal{A}_N],$$

where the $\mathcal{A}_B$ is a subset of size $m \times m$ of $\mathcal{A}$ that is easily invertible. (We can think of $\mathcal{A}$ as an $m \times t(n)$ matrix operating on a vector $\Delta x = \text{vec}(\Delta X)$.) We let

$$\mathcal{Z}(\cdot) := Z\cdot \text{ and } \mathcal{X}(\cdot) := \cdot X.$$

We now divide $\mathcal{Z}, \mathcal{X}, \Delta X$ and ΔZ in the same manner to obtain

$$\Delta X_B = \mathcal{A}_B^{-1} F_p - \mathcal{A}_B^{-1} \mathcal{A}_N \Delta X_N, \tag{3.2}$$

which we substitute back into the dual feasible system (3.1), to obtain the $(n^2 \times t(n))$ system

$$\begin{bmatrix} -\mathcal{Z}_B \mathcal{A}_B^{-1} \mathcal{A}_N + \mathcal{Z}_N & \mathcal{A}^* \mathcal{X} \end{bmatrix} \begin{bmatrix} \Delta X_N \\ \Delta y \end{bmatrix} = \mathcal{X} F_d - F_c - \mathcal{Z}_B \mathcal{A}_B^{-1} F_p. \tag{3.3}$$

Albeit more costly to setup and backsolve than (3.1), this system, for large dense problems is faster to solve because the bottleneck is the least square solve.

We now see that, in contrast to the linear systems used for other directions, the condition number for the RGN linear system stays bounded; thus allowing for more accuracy in the solutions.

Corollary 3.1 *Theorems 2.2 and 2.3 hold with system (2.7) replaced by system (3.3).*

Proof. The proofs follow the same pattern as in the above mentioned theorems. For example, at each iteration we have

$$\mathcal{A}(\Delta X) = 0, \quad -\mathcal{A}^*(\Delta y) + \Delta Z = 0$$

and

$$Z\Delta X + \Delta Z X = 0.$$

Therefore we can apply (2.8) and (2.9).

□

Remark 3.1 *Another approach is to do elimination in the same vein as is done in LP but preserving the least squares properties. This can be done using the explicit structure of the operators in the optimality conditions, see [26, Section 3.2].*

4 NUMERICAL TESTS

4.1 Comparisons Between Directions

We first present some numerical results visually comparing the RGN directions to the well-known Alizadeh-Haeberly-Overton [6] (AHO), Helmberg-Rendl-Vanderbei-Wolkowicz [20]/ Kojima-Shindoh-Hara [24]/ Monteiro [27] (HKM) and Nesterov-Todd [34] (NT) directions. All tests were performed in Matlab using the SDPT3 code [54] from Todd, Toh, Tütüncü. We modified the mainline of SDPT3 to allow further options, namely the GN and the RGN directions. The logic relating to step length and the updating of the centering parameters were untouched to allow a reasonable comparison of the relative value of the GN directions.

The first comparison (Figure 1) illustrates the decrease in both the duality gap and in the infeasibility on random semidefinite programs. The problem has size $(n = 10, m = 15)$ and the algorithm used a predictor-corrector approach on all directions. The RGN direction, seems numerically very stable, and reduces the duality gap as much as the AHO direction in about as many iterations, typically one more. We should note here that the GN direction performs in exactly the same manner as the RGN on those problems, albeit more slowly.

By the RGN direction, we mean the elimination of dy and of part of dX to reduce the system size to $(n^2 \times t(n))$ by forcing the primal and dual feasibility equations. That is, we solved (3.3). On a random SDP, as is the case in this instance, with unknown and dense structure, this elimination will perform well.

The graph shows that the RGN direction performs very well on such problems in terms of the accuracy of the solutions. It does consistently as well or better than AHO.

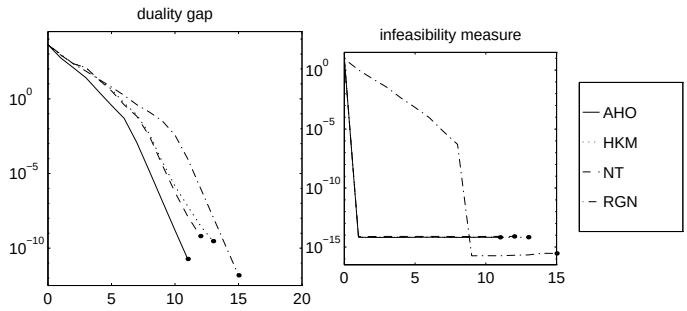

Figure 1: Random Semidefinite Program

The following (Figures 2 and 3) are somewhat unfair tests meant to highlight the robustness of the GN directions with respect to the scaling of the problem. Again, the basic problem is randomly generated with size $(n = 10, m = 15)$, but then the right-hand side vector b is scaled *after the initial point is chosen.* The initial point is therefore badly positioned. The RGN direction is again used.

Remark 4.1 *There are many different RGN directions, i.e. they depend on which basis B is used in (3.3). This basis can be chosen to exploit sparsity or maintain stability. Though the RGN direction was very robust when b was scaled, it did not exhibit any advantage over the other directions when C was scaled. The reason for this is a question for future research.*

Table 1 compares the number of iterations and the accuracy of the solutions of the AHO and RGN directions on random problems. The columns represent, in order, the iteration count, primal feasibility ($||b - \mathcal{A}(X)||/||b||$), dual feasibility ($||\mathcal{A}^*(y) + Z - C||/||C||$), and the duality gap ($\langle X, Z \rangle$). The additional digit of accuracy for the RGN case is most likely due to the presence of a Z^{-1} in the computation of the AHO direction (actually, most semidefinite directions), resulting in a badly conditioned problem close to the solution. The GN systems are always well-conditioned, even close to the solution.

As far as work per iteration is concerned, the major cost of the RGN direction involves the solution of a linear least squares problem of size $n^2 \times t(n)$. This is in contrast to AHO, where the system to solve is of size m, and is then followed by the solution of a Lyapunov equation of order $t(n)$. Therefore, for problems with few constraints, the RGN iterations are more costly and they become competitive as m increases.

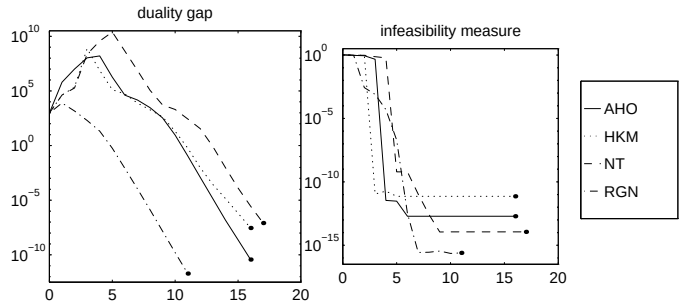

Figure 2: $b \leftarrow b * 100$

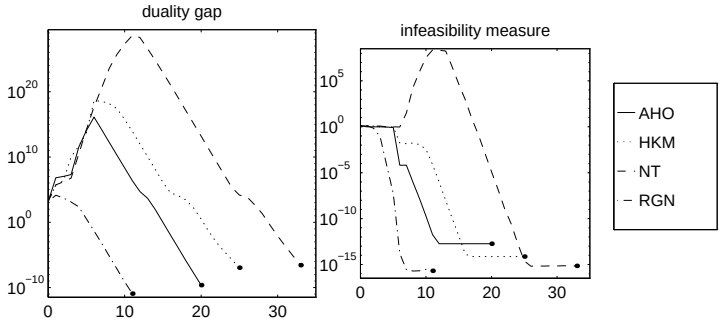

Figure 3: $b \leftarrow b * 500$

AHO				RGN			
Iter	P feas	D feas	Gap	Iter	P feas	D feas	Gap
9	1.2e-14	1.7e-15	4.4e-14	10	2.4e-16	1.8e-15	4.0e-15
9	8.0e-16	2.6e-16	5.8e-14	10	1.2e-16	2.6e-16	4.6e-15
9	1.8e-15	3.7e-16	1.1e-13	10	3.3e-16	5.3e-16	1.0e-14
9	1.6e-15	3.4e-16	7.4e-14	10	5.3e-16	3.9e-16	7.0e-15
9	2.8e-15	5.8e-16	4.5e-14	10	1.4e-16	4.8e-16	2.4e-15
9	1.0e-15	2.8e-16	8.0e-14	10	2.2e-16	1.9e-16	7.3e-15

Table 1: Random problems with $n = 7, m = 28$

AHO				GN		
Size (n)	P feas	D feas	Gap	P feas	D feas	Gap
9	5.7e-04	2.1e-10	5.4e-07	1.8e-06	7.8e-12	6.7e-07
10	7.8e-04	2.8e-10	1.4e-06	4.9e-07	2.2e-12	2.1e-07
11	2.0e-04	1.4e-09	8.6e-07	3.1e-06	3.2e-11	8.9e-07
12	9.2e-04	1.6e-09	2.3e-06	1.5e-06	1.4e-11	5.4e-07
15	6.6e-04	6.0e-09	3.4e-06	3.7e-06	2.8e-11	8.9e-07
AHO				RGN		
20	6.8e-04	1.2e-08	2.9e-05	2.0e-14	6.7e-11	9.7e-07

Table 2: Equipartitioning problems with no Slater point

The final comparison highlights one the main strengths of the GN directions: the numerical accuracy of the solution they provide. The problems under consideration fail to have strictly feasible points in the primal space. That is, the Slater condition fails. The primal constraint force the diagonal elements of the solution to be all ones and the sum of the elements to be zero, making these problems numerically fairly difficult. The primal objective matrix is randomly chosen. Table 2 displays feasibility and the gap of the AHO, GN, and of RGN directions on some of these problems. All stopping criteria were discarded (except for a large maximum iteration count) and the algorithms were left to run to obtain the best possible solution given their respective search directions. It is interesting that the RGN direction attains primal feasibility without losing accuracy in the duality gap.

5 CONCLUSION

In this paper we have studied new search directions for semidefinite programming. The search directions are all based on the GN direction for minimizing the norm of modified, perturbed, optimality conditions. This follows from a new paradigm based on the success in the LP case. Though the paradigm does not provide a new direction for LP, it does for SDP. In fact, the paradigm provides several new directions, and in addition can be used to derive existing directions and show existence where previously existence was a question.

The new approaches have several advantages: e.g. numerical stability, existence and descent generically, scale invariance. And, it has proven to be competitive time wise with existing directions such as AHO and NT and often more robust.

In the future, we still have to study how to calculate the direction more efficiently, how to improve the Gauss-Newton direction using more of the Hessian, and how to exploit

sparsity. In addition, many theoretical questions such as a proof of global convergence remain.

5.1 Newton Versus Gauss-Newton

It is generally accepted that p-d i-p methods obtain their search direction by applying one step of Newton's method to the perturbed optimality conditions. In particular, see e.g. [12], this direction is not obtained from a Newton minimization method on some merit function, since the Jacobian is not symmetric. However, our paradigm points out that it is equivalent to a GN-direction obtained from the least squares minimization problem.

Though the direction from these two points of view are the same, one does get different insights. For example, one might now hope that Newton's method applied to the least squares problem might yield an improved direction. Most textbooks choose the GN-direction over the N-direction only because it is less expensive. In fact, they point out that the error could be large far from the solution and, therefore, a damped Gauss-Newton should be used.

However, the difference between the Gauss-Newton and Newton equations is easily calculated from (2.2). The term dropped in Gauss-Newton involves

$$H := (\Phi_Z F)''(\Phi_Z F) = \begin{bmatrix} 0 & ZX - \mu I \\ ZX - \mu I & 0 \end{bmatrix}$$

This symmetric matrix H has a 0 diagonal and therefore negative eigenvalues off the central path. Therefore, adding this on to the Gauss-Newton equation could cause the direction to point to a maximum or a saddle point.

Many questions arise: can we improve on the GN-direction by properly using the information in H; will a standard or scaled trust region approach mean that the N-direction from the least squares problem is better than the GN-direction; when does the Newton direction point to the minimum and when away from the minimum?

To improve over the GN-direction, one can apply an indefinite scaled trust region approach. Adding the indefinite trust region constraint

$$\langle H\Delta s, \Delta s\rangle = \delta,$$

for an appropriate δ results in the same direction as the GN-direction, since the H will cancel off. Therefore, we can do at least as well as the GN-direction by being clever enough. To improve on this, we need to either find a better scaling than H or adaptively choose δ to get closer to the boundary and improve the merit (least squares) function.

References

[1] A. ALFAKIH, A. KHANDANI, and H. WOLKOWICZ. An interior-point method for approximate distance matrix completions. Technical Report CORR Report 98-9, University of Waterloo, 1997. URL: ftp://orion.uwaterloo.ca/pub/henry/reports/distmat.ps.gz.

[2] F. ALIZADEH. *Combinatorial optimization with interior point methods and semidefinite matrices.* PhD thesis, University of Minnesota, 1991.

[3] F. ALIZADEH. Interior point methods in semidefinite programming with applications to combinatorial optimization. *SIAM Journal on Optimization*, 5:13–51, 1995.

[4] F. ALIZADEH, J.-P. HAEBERLY, M. V. NAYAKKANKUPPAM, and M. L. OVERTON. Sdppack user's guide – version 0.8 beta. Technical Report TR1997–734, Courant Institute of Mathematical Sciences, NYU, New York, NY, March 1997.

[5] F. ALIZADEH, J-P.A. HAEBERLY, and M.L. OVERTON. Complementarity and nondegeneracy in semidefinite programming. *Mathematical Programming*, 77:111–128, 1997.

[6] F. ALIZADEH, J-P.A. HAEBERLY, and M.L. OVERTON. Primal-dual interior-point methods for semidefinite programming: convergence rates, stability and numerical results. *Siam J. Optimization*, 1998. to appear in SIOPT.

[7] ADI BEN-ISRAEL. Newton's method with modified functions. In *Recent developments in optimization theory and nonlinear analysis (Jerusalem, 1995)*, volume 204 of *Contemp. Math.*, pages 39–50. Amer. Math. Soc., Providence, RI, 1997.

[8] A. BJORCK. *Numerical Methods for Least Squares Problems.* SIAM, Philadelphia, 1996.

[9] S. BOYD, L. El GHAOUI, E. FERON, and V. BALAKRISHNAN. *Linear Matrix Inequalities in System and Control Theory*, volume 15 of *Studies in Applied Mathematics.* SIAM, Philadelphia, PA, June 1994.

[10] N. BRIXIUS, F.A. POTRA, and R. SHENG. SDPHA: a matlab implementation of homogeneous interior-point algorithms for semidefinite programming. Technical report, University of Iowa, Iowa City, IA, 1998.

[11] Jr. DENNIS, J. E. and ROBERT B. SCHNABEL. *Numerical methods for unconstrained optimization and nonlinear equations*, volume 16 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1996. Corrected reprint of the 1983 original.

[12] A.S. EL-BAKRY, R.A. TAPIA, T. TSUCHIYA, and Y. ZHANG. On the formulation and theory of the Newton interior-point method for nonlinear programming. *Journal of Optimization Theory and Applications*, 89:507–541, 1996.

[13] M. ELHALLABI. An inexact hybrid algorithm for nonlinear systems of equations. Technical Report CRPC-TR94372, Department of Mathematical Sciences, Rice University, Houston, TX, 1994.

[14] M. ELHALLABI and R.A. TAPIA. An inexact trust-region feasible-point algorithm: globalization of Gauss-Newton method. Technical Report TR93-43, Department of Mathematical Sciences, Rice University, Houston, TX, 1993.

[15] M. ELHALLABI and R.A. TAPIA. An inexact trust-region feasible-point algorithm for nonlinear systems of equalities and inequalities. Technical Report CRPC-TR95543, Department of Mathematical Sciences, Rice University, Houston, TX, 1995.

[16] K. FUJISAWA, M. KOJIMA, and K. NAKATA. SDPA semidefinite programming algorithm. Technical report, Dept. of Information Sciences, Tokyo Institute of Technology, Tokyo, Japan, 1995.

[17] J. GERLACH. Accelerated convergence in Newton's method. *SIAM Review*, 36:272–276, 1994.

[18] M. X. GOEMANS. Semidefinite programming in combinatorial optimization. *Mathematical Programming*, 79:143–162, 1997.

[19] C. HELMBERG. *An interior point method for semidefinite programming and max-cut bounds*. PhD thesis, Graz University of Technology, Austria, 1994.

[20] C. HELMBERG, F. RENDL, R. J. VANDERBEI, and H. WOLKOWICZ. An interior point method for semidefinite programming. *SIAM Journal on Optimization*, pages 342–361, 1996.
URL: ftp://orion.uwaterloo.ca/pub/henry/reports/sdp.ps.gz.

[21] C. JOHNSON, B. KROSCHEL, and H. WOLKOWICZ. An interior-point method for approximate positive semidefinite completions. *Computational Optimization and Applications*, 9(2):175–190, 1998.

[22] S. KARISCH. *Nonlinear Approaches for Quadratic Assignment and Graph Partition Problems.* PhD thesis, University of Graz, Graz, Austria, 1995.

[23] E. DE KLERK. *Interior point methods for semidefinite programming.* PhD thesis, Delft University, 1997.

[24] M. KOJIMA, S. SHINDOH, and S. HARA. Interior-point methods for the monotone semidefinite linear complementarity problem in symmetric matrices. *SIAM J. Optim.*, 7(1):86–125, 1997.

[25] S. KRUK. Semidefinite programming applied to nonlinear programming. Master's thesis, University of Waterloo, 1996.
URL: http://orion.uwaterloo.ca/~hwolkowi/henry/software/sqqp.d/krukmsthesis.ps.gz.

[26] S. KRUK, M. MURAMATSU, F. RENDL, R.J. VANDERBEI, and H. WOLKOWICZ. The Gauss-Newton direction in linear and semidefinite programming. Technical Report CORR 98-16, University of Waterloo, Waterloo, Canada, 1998. Detailed Web Version at
URL: ftp://orion.uwaterloo.ca/pub/henry/reports/gnsdplong.ps.gz.

[27] R.D.C. MONTEIRO. Primal-dual path-following algorithms for semidefinite programming. *SIAM J. Optim.*, 7(3):663–678, 1997.

[28] R.D.C. MONTEIRO and T. TSUCHIYA. Polynomial convergence of a new family of primal-dual algorithms for semidefinite programming. Technical Report No. 627, The Institute of Statistical Mathematics, Tokyo, Japan, 1997.

[29] R.D.C. MONTEIRO and Y. ZANJACOMO. A note on the existence of the Alizadeh-Haeberly-Overton direction for semidefinite programming. *Mathematical Programming*, 78:393–401, 1997.
URL http://www.isye.gatech.edu/~monteiro/tech_reports/note.ps.

[30] R.D.C. MONTEIRO and Y. ZHANG. A unified analysis for a class of path-following primal-dual interior-point algorithms for semidefinite programming. Technical report, Georgia Tech, Atlanta, GA, 1996. ftp://pc5.math.umbc.edu/pub/mz_sdp.ps.gz or ftp://pc5.math.umbc.edu/pub/mz_sdp.dvi.gz.

[31] J.J. MORÉ and D.C. SORENSEN. Computing a trust region step. *SIAM J. Sci. Statist. Comput.*, 4:553–572, 1983.

[32] L. NAZARETH. Computer solution of linear programs: non-simplex algorithms II. Technical report, Dept. of Pure and Applied Mathematics, Washington State University, Pullman, 1996.

[33] Y. E. NESTEROV and A. S. NEMIROVSKI. *Interior Point Polynomial Algorithms in Convex Programming : Theory and Algorithms.* SIAM Publications. SIAM, Philadelphia, USA, 1994.

[34] Y. E. NESTEROV and M. TODD. Primal-dual interior-point methods for self-scaled cones. Technical Report TR1125, School of OR and IE, Cornell University, Ithaca, NY, 1995.

[35] P. PARDALOS and H. WOLKOWICZ (Editors). *Topics in Semidefinite and Interior-Point Methods*, volume 18 of *The Fields Institute for Research in Mathematical Sciences, Communications Series.* American Mathematical Society, Providence, Rhode Island, 1998.

[36] G. PATAKI. *Cone programming and eigenvalue optimization: geometry and algorithms.* PhD thesis, Carnegie Mellon University, Pittsburgh, PA, 1996.

[37] G. PATAKI and L. TUNCEL. On the generic properties of convex optimization problems in conic form. Technical Report CORR 97-16, Department of Combinatorics and Optimization, Waterloo, Ont, 1997.

[38] M.J.D. POWELL. The use of band matrices for second derivative approximations in trust region algorithms. DAMTP Report 1997/NA12, Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, England, 1997.

[39] M. RAMANA, L. TUNCEL, and H. WOLKOWICZ. Strong duality for semidefinite programming. *SIAM Journal on Optimization*, 7(3):641–662, 1997. URL:ftp://orion.uwaterloo.ca/pub/henry/reports/strongdual.ps.gz.

[40] M.V. RAMANA. *An algorithmic analysis of multiquadratic and semidefinite programming problems.* PhD thesis, Johns Hopkins University, Baltimore, Md, 1993.

[41] F. RENDL. Semidefinite programming and combinatorial optimization. Technical report, University of Graz, Graz, Austria, 1997.

[42] F. RENDL, R. J. VANDERBEI, and H. WOLKOWICZ. Max-min eigenvalue problems, primal-dual interior point algorithms, and trust region subproblems. *Optimization Methods and Software*, 5:1–16, 1995.

[43] F. RENDL and H. WOLKOWICZ. A semidefinite framework for trust region subproblems with applications to large scale minimization. *Mathematical Programming*, 77(2):273–299, 1997.

[44] R. SAIGAL, L. VANDENBERGHE, and H. WOLKOWICZ (Editors). *Handbook on: SEMIDEFINITE PROGRAMMING AND APPLICATIONS*. Kluwer Academic Publishers, Waterloo, Canada, 1998. in progress.

[45] A. SHAPIRO. First and second order analysis of nonlinear semidefinite programs. *Mathematical Programming*, 77:301–320, 1997.

[46] A. SHAPIRO and M.K.H. FAN. On eigenvalue optimization. *SIAM J. Opt.*, 5:552–569, 1995.

[47] R. SHENG and F.A. POTRA. Nonsymmetric search directions for semidefinite programming. Technical report, University of Iowa, Iowa City, IA, 1997.

[48] M. SHIDA and S. SHINDOH andM. KOJIMA. Existence of search directions in interior-point algorithms for the SDP and monotone SDLCP. Technical report, Dept. of Information Sciences, Tokyo Institute of Technology, Tokyo, Japan, 1996.

[49] D.C. SORENSEN. Newton's method with a model trust region modification. *SIAM J. Numer. Anal.*, 19:409–426, 1982.

[50] J. STURM. *Primal-dual interior point approach to semidefinite programming*. PhD thesis, Erasmus University Rotterdam, 1997.

[51] C. SUN. Dealing with dense rows in the solution of sparse linear least squares problems. Technical Report TR95-227, Department of Computer Science, Cornell University, Ithaca, NY, 1995.

[52] (Editor) T. TERLAKY. *Interior Point Methods of Mathematical Programming*. Kluwer Academic Publishers, Dordrecht; Boston, 1996.

[53] M. TODD. On search directions in interior-point methods programming. Technical Report TR1205, School of OR and IE, Cornell University, Ithaca, NY, 1997.

[54] M. TODD, K.C. TOH, and R.H. TUTUNCU. A MATLAB software package for semidefinite programming. Technical report, School of OR and IE, Cornell University, Ithaca, NY, 1996.

[55] M. TODD, K.C. TOH, and R.H. TUTUNCU. On the nesterov-todd direction in semidefinite programming. Technical Report TR1154, School of OR and IE, Cornell University, Ithaca, NY, 1996.

[56] L. VANDENBERGHE and S. BOYD. Semidefinite programming. *SIAM Review*, 38:49–95, 1996.

[57] S. WRIGHT. *Primal-Dual Interior-Point Methods.* SIAM, Philadelpia, Pa, 1996.

[58] Y. YE. *Interior Point Algorithms: Theory and Analysis.* Wiley-Interscience series in Discrete Mathematics and Optimization. John Wiley & Sons, New York, 1997.

[59] Y. ZHANG. Basic equalities and inequalities in primal-dual interior-point methods for semidefinite programming. Technical report, Department of Mathematics and Statistics, University of Maryland Baltimore County, 1995.

[60] Q. ZHAO. *Semidefinite Programming for Assignment and Partitioning Problems.* PhD thesis, University of Waterloo, 1996.
URL: ftp://orion.uwaterloo.ca/pub/henry/software/qap.d/zhaophdthesis.ps.gz.

Note that a comprehensive bibliography is available over WWW with URL:
http://orion.uwaterloo.ca/~hwolkowi/henry/reports/psd.bib.gz