A primal-dual potential reduction method for problems involving matrix inequalities

Lieven Vandenberghe[1]

Katholieke Universiteit Leuven
Departement Elektrotechniek, Afdeling ESAT
K. Mercierlaan 94, 3001 Leuven, Belgium
Internet: `vandenbe@isl.stanford.edu`

Stephen Boyd[2]

Information Systems Laboratory
Electrical Engineering Department
Stanford University
Stanford CA 94305
Internet: `boyd@isl.stanford.edu`

May 1994

[1]Senior research assistant of the Belgian National Fund for Scientific Research (NFWO). Research supported in part by the Belgian program on interuniversity attraction poles (IUAP 17 and 50) initiated by the Belgian State, Prime Minister's Office, Science Policy Programming.

[2]Research supported in part by AFOSR (under F49620-92-J-0013), NSF (under ECS-9222391), and ARPA (under F49620-93-1-0085).

Abstract

We describe a potential reduction method for convex optimization problems involving matrix inequalities. The method is based on the theory developed by Nesterov and Nemirovsky and generalizes Gonzaga and Todd's method for linear programming. A worst-case analysis shows that the number of iterations grows as the square root of the problem size, but in practice it appears to grow more slowly. As in other interior-point methods the overall computational effort is therefore dominated by the least-squares system that must be solved in each iteration. A type of conjugate-gradient algorithm can be used for this purpose, which results in important savings for two reasons. First, it allows us to take advantage of the special structure the problems often have (*e.g.*, Lyapunov or algebraic Riccati inequalities). Second, we show that the polynomial bound on the number of iterations remains valid even if the conjugate-gradient algorithm is not run until completion, which in practice can greatly reduce the computational effort per iteration.

We describe in detail how the algorithm works for optimization problems with L Lyapunov inequalities, each of size m. We prove an overall *worst-case* operation count of $O(m^{5.5}L^{1.5})$. The *average-case* complexity appears to be closer to $O(m^4L^{1.5})$. This estimate is justified by extensive numerical experimentation, and is consistent with other researchers' experience with the practical performance of interior-point algorithms for linear programming.

This result means the computational cost of extending current control theory that is based on the solution of Lyapunov or Riccati *equations* to a theory that is based on the solution of (multiple, coupled) Lyapunov or Riccati *inequalities* is modest.

Key words: interior point algorithms, linear matrix inequalities, positive-definite programming.

Abbreviated title: Primal-dual method for matrix inequalities.

To be published in *Mathematical Programming Series B,* special issue on nondifferentiable and large scale optimization, Revised June 1993 and May 1994.

1 Introduction

1.1 Motivation

Many problems in systems and control theory can be formulated (or reformulated) as optimization problems involving *linear matrix inequalities*, *i.e.*, constraints requiring an affine combination of symmetric matrices to be positive semidefinite. Reference [9] gives a broad survey of such problems.

These matrix inequalities are usually highly structured. One typical example is the (convex) *Lyapunov inequality* which has the form

$$APB + B^T PA^T + D \geq 0,$$

where the square matrices A, B and D are given, D is symmetric, and the symmetric matrix P is the optimization variable. Another important example is the (convex) *algebraic Riccati inequality:*

$$A^T P + PA + PBR^{-1} B^T P + Q \leq 0,$$

where A, B, Q and R are given, Q is symmetric, R is positive definite, and the matrix P is the optimization variable. This quadratic matrix inequality can be recast as a linear matrix inequality which is very similar in form to the Lyapunov inequality:

$$\begin{bmatrix} -A^T P - PA - Q & PB \\ B^T P & R \end{bmatrix} = \bar{A} P \bar{B} + \bar{B}^T P \bar{A}^T + \bar{D} \geq 0,$$

where

$$\begin{aligned} \bar{A} &= \begin{bmatrix} I \\ 0 \end{bmatrix}, \\ \bar{B} &= \begin{bmatrix} -A & B \end{bmatrix}, \\ \bar{D} &= \begin{bmatrix} -Q & 0 \\ 0 & R \end{bmatrix}. \end{aligned}$$

Lyapunov and Riccati inequalities arise, for example, in stability analysis of dynamical systems.

1.2 A typical problem

In this paper we describe a potential reduction method for general convex optimization problems involving matrix inequalities such as Lyapunov or Riccati inequalities. We give complete details for the typical problem of minimizing a linear functional of an $m \times m$ matrix P subject to L Lyapunov inequalities:

$$\min \mathbf{Tr}\, EP, \quad A_k P B_k + B_k^T P A_k^T + D_k \geq 0, \quad k = 1, \ldots, L, \tag{1}$$

where E, A_k, B_k, D_k, $k = 1, \ldots, L$, are $m \times m$ matrices with E and D_k symmetric.

We will show that, in the *worst case*, the algorithm we describe takes $O(m^{5.5}L^{1.5})$ operations to solve the optimization problem (1). The *average case* complexity appears to increase much more slowly with m, as $O(m^{\beta}L^{\gamma})$, with $\beta \approx 4$ and $\gamma \approx 1.5$. To appreciate these numbers, consider the following. A single Lyapunov equation $APB + B^TPA^T + D = 0$ (which is just a set of $m(m+1)/2$ linear equations for the $m(m+1)/2$ variables in P) can be solved in $O(m^3)$ operations by exploiting the special structure of the equations (see, *e.g.*, [20]). Therefore, it takes $O(m^3L)$ operations to solve L independent Lyapunov equations. Comparing this operation count to $O(m^4L^{1.5})$, we see that the relative cost of solving L coupled Lyapunov inequalities, compared to solving L independent Lyapunov equations, is only a factor of $mL^{0.5}$. A similar statement holds for Riccati inequalities.

Much of modern control theory involves the solution of Riccati and Lyapunov equations. Our results show that the computational cost of extending current control theory to a theory based on the solution of (multiple, coupled) Lyapunov or Riccati inequalities is modest. (Extensive discussion of this topic can be found in [9].)

We also note that the problem (1) includes linear programming as a special case. When E, A_k, B_k, D_k, $k = 1, \ldots, L$ are all diagonal matrices, the problem reduces to minimizing a linear function subject to a set of linear inequalities.

1.3 A brief historical overview

A fairly complete history of matrix inequalities arising in control theory can be found in [9]. Problems of this type also occur in statistics [35, 12, 13] and structural analysis and design [7, 31, 32]. In a very early paper, Bellman and Ky Fan [6] discuss matrix inequalities from an optimization viewpoint, and describe optimality conditions, duality, and theorems of alternatives.

Interest in interior-point polynomial time methods for more general nonlinear convex problems, and problems involving matrix inequalities in particular, started soon after the publication of interior-point methods for linear programming. Initial efforts were directed towards generalizing the method of centers and other central path-following methods [30, 8, 10, 24]. More recently, potential reduction methods have been extended to problems involving matrix inequalities. Nesterov and Nemirovsky [30, Ch.4] describe three potential reduction algorithms for problems involving matrix inequalities: a generalization of Karmarkar's method, a projective method, and a generalization of the method of Ye [36]. In [3, 2],Alizadeh describes several potential reduction methods for problems involving matrix inequalities, emphasizing their similarity to the analogous methods for linear programming. These potential reduction methods all share an important advantage over the earlier path-following methods: they allow so-called "large steps," *i.e.*, the use of (computationally cheap) line search procedures to reduce the potential function at each iteration by an amount that is often substantially more than is guaranteed by the complexity analysis.

The algorithm that we describe in this paper involves two important extensions beyond the methods described by Nesterov, Nemirovsky, and Alizadeh. First, it takes advantage of the special structure of the matrix inequalities we encounter, *e.g.*, Lyapunov or Riccati. Second, it allows the use of approximate search directions, which can be computed by a conjugate gradient algorithm.

The algorithm can also be considered as extending Gonzaga and Todd's primal-dual algorithm for linear programming [22] in two ways: to handle problems involving linear matrix inequalities, and to use approximate search directions. The use of approximate search directions is related to the "relaxed version" of Karmarkar's method for linear programming, as described by Goldfarb and Mehrotra [17, 18].

1.4 Outline

In section 2 we describe the problem and the algorithm in general terms, deferring proofs and further details to the appendix or later sections. In section 3 we give the details of computing a suitable search direction via an appropriate conjugate gradient method. This is the key to exploiting the special structure of the matrix inequalities as well as reducing the computational effort per iteration by using approximate search directions. In section 4 we describe an efficient plane search technique for computing step lengths at each iteration.

In section 5 we return to the specific problem (1) of minimizing a linear function of a matrix subject to a set of Lyapunov inequalities, giving details of the conjugate gradient algorithm. In section 6 we present the results of some numerical experiments that support our claim that the average number of operations needed to solve problems like (1) increases only as $O(m^4 L^{1.5})$.

1.5 Notation

- The integer n determines the size of the general problem. When it is not relevant or not varying, we will suppress the dependence on n in notation.
- $\mathcal{S}$ will denote the space of symmetric $n \times n$ matrices. When we need to be specific about n, we will write it as $\mathcal{S}_n$.
- $\mathcal{P}$ will denote the cone of symmetric $n \times n$ positive-semidefinite matrices.
- I will denote the identity matrix, with size determined from context.
- The inner product of two matrices in $\mathcal{S}$ is defined as $\langle X, Z \rangle = \mathbf{Tr}\,(XZ)$. The corresponding norm is the Frobenius norm, denoted $\|X\| = (\mathbf{Tr}\,X^2)^{1/2}$.
- $X^{1/2}$ denotes the symmetric square-root of $X \in \mathcal{P}$.
- The direct sum of two matrices X and Z is written as

$$X \oplus Z = \begin{bmatrix} X & 0 \\ 0 & Z \end{bmatrix}$$

Similarly, if $\mathcal{L}^1$ and $\mathcal{L}^2$ are subspaces in $\mathcal{S}$ we write

$$\mathcal{L}^1 \oplus \mathcal{L}^2 = \{X \oplus Z \mid X \in \mathcal{L}^1, Z \in \mathcal{L}^2\} \subseteq S_{2n}.$$

– The orthogonal complement of a subspace $\mathcal{L} \subseteq \mathcal{S}$ will be denoted $\mathcal{L}^\perp$, *i.e.*, $\mathcal{L}^\perp = \{Z \in \mathcal{S} | \langle X, Z\rangle = 0 \text{ for all } X \in \mathcal{L}\}$.

2 Algorithm: general formulation

2.1 Conic formulation

We will express the general problem using the conic formulation of Nesterov and Nemirovsky [30, §4.1]:

$$\min \langle C, X\rangle, \quad X \in \mathcal{P} \cap (\mathcal{L} + D). \tag{2}$$

Here C, D are given elements in $\mathcal{S}$ and $\mathcal{L}$ is a subspace of $\mathcal{S}$. The associated dual problem is:

$$\min \langle Z, D\rangle, \quad Z \in \mathcal{P} \cap (\mathcal{L}^\perp + C). \tag{3}$$

Matrices X and Z will be called (primal and dual) *feasible* if they belong to $\mathcal{P} \cap (\mathcal{L} + D)$ and $\mathcal{P} \cap (\mathcal{L}^\perp + C)$, and *strictly feasible* if in addition they lie in **Int** $\mathcal{P}$. For a pair of feasible matrices X, Z, the quantity $\langle X, Z\rangle$ is the *duality gap* for the primal and dual problems. We note the following facts:

– Since X, $Z \in \mathcal{P}$, $\langle X, Z\rangle \geq 0$ and $\langle X, Z\rangle = 0$ if and only if $XZ = 0$.

– The duality gap is actually affine in X, Z, and not a bilinear form as it appears at first sight:

$$X \in \mathcal{L} + D, \;\; Z \in \mathcal{L}^\perp + C \quad \Longrightarrow \quad \langle X, Z\rangle = \langle C, X\rangle + \langle Z, D\rangle - \langle C, D\rangle. \tag{4}$$

– $\langle X, Z\rangle$ is an upper bound on the difference between the value of the primal objective with X and the optimal value of the primal problem, *i.e.*,

$$\langle C, X\rangle - \inf\{\langle C, \tilde{X}\rangle | \tilde{X} \in \mathcal{P} \cap (\mathcal{L} + D)\} \leq \langle X, Z\rangle.$$

Similarly,

$$\langle Z, D\rangle - \inf\{\langle \tilde{Z}, D\rangle | \tilde{Z} \in \mathcal{P} \cap (\mathcal{L}^\perp + C)\} \leq \langle X, Z\rangle.$$

(These follow immediately from the two preceding observations.)

We make the following assumption about the pair of problems (2) and (3):

Assumption 1 *We are given strictly feasible primal and dual matrices* $X^{(0)}$ *and* $Z^{(0)}$.

This is precisely the assumption made by Nesterov and Nemirovsky in their potential reduction algorithms [30, Ch.4]. It has the following implications [30, §4.2]:

- The feasible sets $\mathcal{P} \cap (\mathcal{L} + D)$ and $\mathcal{P} \cap (\mathcal{L}^\perp + C)$ have nonempty relative interiors.
- The primal and dual objective functions are bounded below on the feasible sets.
- The primal and dual problems are solvable. X solves (2) if and only if there exists a dual feasible Z with $\langle X, Z\rangle = 0$. Similarly, Z solves (3) if and only if there exists a primal feasible X with $\langle X, Z\rangle = 0$.

Several methods are known to circumvent this assumption. The easiest is to precede the algorithm with a phase I algorithm to find feasible initial points (see, *e.g.*, [30, §4.3.5]). In other approaches both phases are combined; see, *e.g.*, [28, 26, 14, 4].

Problems (2) and (3) together are therefore equivalent to a linear optimization problem with known optimal value zero:

$$\min\langle X, Z\rangle, \quad X \in \mathcal{P} \cap (\mathcal{L} + D), \quad Z \in \mathcal{P} \cap (\mathcal{L}^\perp + C).$$

2.2 Central path

The material of this section is not needed either in the description of the algorithm or in the proofs. We present it because it allows us to give important interpretations and useful insight.

For $X \in \textbf{Int}\ \mathcal{P}$ we define

$$F(X) = \log\det X^{-1}.$$

F is strictly convex and converges to ∞ as X approaches the boundary of $\mathcal{P}$, *i.e.*, F is a barrier function for the positive definite cone $\textbf{Int}\ \mathcal{P}$. This function has very simple first and second derivatives. The gradient at a point $X \in \textbf{Int}\ \mathcal{P}$ is

$$\nabla F(X) = -X^{-1}.$$

The Hessian $\nabla^2 F(X)$, when considered as a mapping from $\mathcal{S}$ to $\mathcal{S}$, is given by a congruence operation: for $H \in \mathcal{S}$,

$$\nabla^2 F(X) H = X^{-1} H X^{-1}. \tag{5}$$

For $\alpha > 0$ consider the set of strictly feasible pairs X, Z with $\langle X, Z\rangle = \alpha$. Since $\langle X, Z\rangle = \alpha$ is an affine constraint on X, Z, this set is the intersection of $\textbf{Int}\ \mathcal{P} \oplus \textbf{Int}\ \mathcal{P}$ with an affine set. It can be shown that under Assumption 1 this set is nonempty and bounded. The *analytic center* of this set is the minimizer of $F(X) + F(Z)$, or equivalently, the matrix with maximum determinant. We denote the analytic center as $X^*(\alpha), Z^*(\alpha)$:

$$(X^*(\alpha), Z^*(\alpha)) = \mathop{\mathrm{argmin}}_{\substack{X \in \mathcal{P} \cap (\mathcal{L} + D) \\ Z \in \mathcal{P} \cap (\mathcal{L}^\perp + C) \\ \langle X, Z\rangle = \alpha}} F(X) + F(Z). \tag{6}$$

(Since the feasible set here is bounded, $F(X)+F(Z)$ is bounded below.) Thus, among all feasible pairs X, Z with the duality gap α, the pair X^*, Z^* maximizes $\det(XZ)$. Roughly speaking, we can consider (X^*, Z^*) as the pair with duality gap α that is "most feasible".

The curve given by (X^*, Z^*) for $\alpha > 0$ is called the *path of centers* for the problems (2) and (3). Evidently (X^*, Z^*) converge to a primal and dual optimal pair as $\alpha \to 0$. The central pair (X^*, Z^*) has many important properties. For our purposes here we need:

Theorem 1 *$X^*(\alpha)Z^*(\alpha) = (\alpha/n)I$. Conversely, if X and Z are a feasible pair and $XZ = (\alpha/n)I$ then $X = X^*(\alpha)$ and $Z = Z^*(\alpha)$.*

In other words, centrality is characterized by X and Z being inverses of each other, up to a constant. The proof is given in the appendix.

Now consider a feasible pair (X, Z), and define $\alpha = \langle X, Z\rangle$. Then $(X^*(\alpha), Z^*(\alpha))$ is the central pair with the same duality gap as X, Z. Therefore

$$F(X) + F(Z) \geq F(X^*(\alpha)) + F(Z^*(\alpha)) = -\log\det(X^*Z^*) = n\log n - n\log\langle X, Z\rangle$$

with equality holding only when X, Z are central. We will see that the difference between $F(X)+F(Z)$ and $F(X^*(\alpha)) + F(Z^*(\alpha))$ can be interpreted as a measure of the deviation of (X, Z) from centrality. This difference is:

$$\psi(X, Z) = F(X) + F(Z) + n\log\langle X, Z\rangle - n\log n.$$

We have already observed that $\psi(X, Z) \geq 0$ for all feasible X, Z and $\psi(X, Z) = 0$ only if X, Z is on the central path. Moreover, ψ approaches ∞ as the pair (X, Z) approaches the boundary of $\mathcal{P} \oplus \mathcal{P}$. Therefore ψ satisfies some basic requirements for a measure of deviation from centrality. We will see below a much more specific interpretation of ψ as a measure of deviation from centrality.

We note that ψ is not convex or quasiconvex (except of course when restricted to $\langle X, Z\rangle$ constant). We also note that ψ depends only on the eigenvalues $\lambda_1, \ldots, \lambda_n$ of XZ:

$$\psi(X, Z) = n\log\frac{\left(\sum_{i=1}^n \lambda_i\right)/n}{\left(\prod_{i=1}^n \lambda_i\right)^{1/n}}.$$

Thus $\psi(X, Z)$ is n times the logarithm of the ratio of the arithmetic to geometric mean of the eigenvalues of XZ. (From which we see again that ψ is nonnegative, and zero only when XZ is a multiple of the identity.) We can also think of ψ as a smooth measure of condition number of the matrix XZ since

$$\log\kappa - 2\log 2 \leq \psi(X, Z) \leq (n-1)\log\kappa$$

where $\kappa = \lambda_{\max}/\lambda_{\min}$ is the condition number of XZ.

We can give a nice interpretation of ψ using Nesterov and Nemirovsky's theory. We consider the problem of computing X^*, Z^* given X, Z. In [30, §2.2] Nesterov and Nemirovsky give a very simple damped Newton algorithm for computing X^*, Z^* that has the following properties:

– Until the region of quadratic convergence is reached, the objective $F(X)+F(Z)$ decreases at least by the absolute constant 0.3068 at each Newton step. (By absolute constant we mean it does not depend on n, the problem data, or the required accuracy of computing X^*, Z^*.)

– Once the region of quadratic convergence is reached, at most a *constant* number c of Newton steps is required to compute X^*, Z^* to a given accuracy. (The constant c does not depend on n or the problem data, but only on the required accuracy ϵ. Since the convergence is quadratic in this region, c grows as $\log\log 1/\epsilon$ if ϵ decreases.)

As a consequence we see that the number of Newton steps required to compute X^*, Z^* given X, Z can be bounded in terms of $F(X)+F(Z)-F(X^*)-F(Z^*)=\psi(X,Z)$:

$$\#\text{Newton steps} \quad \leq \quad c + 3.26\psi(X,Z) \tag{7}$$

where c depends only on the required accuracy of computing X^*, Z^* and grows extremely slowly. In other words: $\psi(X,Z)$ is, up to a constant, an upper bound on the computational effort required to "center" (X,Z) (meaning, compute the central pair with the same duality gap).

2.3 Potential function

Let $\nu \geq 1$. For strictly feasible X and Z, we define the primal-dual potential function as

$$\begin{aligned} \phi(X,Z) &= \nu\sqrt{n}\log\langle X,Z\rangle + \psi(X,Z) && (8)\\ &= (n+\nu\sqrt{n})\log\langle X,Z\rangle + F(X)+F(Z) - n\log n && (9)\\ &= (n+\nu\sqrt{n})\log\sum_{i=1}^{n}\lambda_i - \sum_{i=1}^{n}\log\lambda_i - n\log n && (10)\end{aligned}$$

where λ_i are the eigenvalues of XZ.

We can interpret the two terms in (8) as follows. The first term, $\nu\sqrt{n}\log\langle X,Z\rangle$, depends only on the duality gap, and decreases to $-\infty$ as the duality gap approaches zero. Therefore, a fixed decrease in the first term corresponds to a fixed fractional reduction of the duality gap.

We have already noted that the second term, $\psi(X,Z)$, can be interpreted as a measure of deviation from centrality of the pair X Z, and increases from zero on the central path to $+\infty$ as X or Z approach the boundary of $\mathcal{P}$. A fixed decrease in the second term corresponds to a fixed amount of "centering" in the following sense: up to a constant, it is the reduction in the (bound on) computational effort required to "center" the current pair. Note that along the central path, $\phi(X^*(\alpha), Z^*(\alpha)) = \nu\sqrt{n}\log\alpha$ which decreases to $-\infty$ as α converges to zero.

The constant ν evidently determines the relative weight of the two terms, which measure the duality gap and the deviation from centrality. Using our interpretation (7) we can be more specific. If ψ decreases by one, the new pair is 3.26 Newton steps closer to centrality than the original pair (or more precisely, the upper bound on the number of Newton steps required to center the new pair is 3.26 smaller than the upper bound for the original pair). If the other term, $\nu\sqrt{n}\log\langle X,Z\rangle$, decreases by one, then the duality gap is reduced by the factor $\exp(-1/\nu\sqrt{n}) \approx 1 - 1/(\nu\sqrt{n})$. In other words,

$$\nu\sqrt{n} \text{ fewer Newton steps to center } \sim \text{ 31\% duality gap reduction}$$

where $\sim$ means that the left and right-hand sides result in an equal decrease in ϕ.

By minimizing the smooth function ϕ, we solve the primal and dual problems (2) and (3). Indeed since $\psi(X,Z) \geq 0$ for feasible X and Z, we have

$$\langle X,Z\rangle \leq \exp\frac{\phi(X,Z)}{\nu\sqrt{n}} \tag{11}$$

which shows that small potential implies small duality gap.

For notational convenience we will often use $W = X \oplus Z$, and $\phi(W)$ to denote $\phi(X,Z)$. For future reference we give the gradient of ϕ:

$$\nabla\phi(X,Z) = \left(\frac{n+\nu\sqrt{n}}{\langle X,Z\rangle}C - X^{-1}\right) \oplus \left(\frac{n+\nu\sqrt{n}}{\langle X,Z\rangle}D - Z^{-1}\right) = \rho(C \oplus D) - W^{-1}$$

where we define ρ as

$$\rho = \frac{n+\nu\sqrt{n}}{\langle X,Z\rangle}. \tag{12}$$

2.4 Algorithm

The basic idea of the primal-dual algorithm is to generate iterations of primal and dual feasible matrices satisfying

$$\phi(X^{(k+1)}, Z^{(k+1)}) \leq \phi(X^{(k)}, Z^{(k)}) - \delta, \tag{13}$$

for some absolute positive constant δ. By (11) we therefore have:

$$\langle X^{(k)}, Z^{(k)}\rangle \leq \exp\frac{\phi(X^{(0)}, Z^{(0)}) - k\delta}{\nu\sqrt{n}} = \langle X^{(0)}, Z^{(0)}\rangle c_0 c_1^k \tag{14}$$

where

$$c_0 = \exp\frac{\psi(X^{(0)}, Z^{(0)})}{\nu\sqrt{n}}, \qquad c_1 = \exp\frac{-\delta}{\nu\sqrt{n}}.$$

We can interpret the result (14) as follows: the duality gap converges to zero at least exponentially at a rate given by the constant c_1. The constant c_0 depends only on the centrality of the initial pair, and is one if the initial pair is central.

In other words, we have polynomial convergence:

Theorem 2 *Assume that (13) holds with some $\delta > 0$ that does not depend on n or ϵ, where $0 < \epsilon < 1$. Then for*

$$k \geq \frac{\nu\sqrt{n}\log(1/\epsilon) + \psi(X^{(0)}, Z^{(0)})}{\delta}$$

we have $\langle X^{(k)}, Z^{(k)}\rangle < \epsilon\langle X^{(0)}, Z^{(0)}\rangle$.

Roughly speaking, we have convergence in $O(\sqrt{n})$ steps, provided the initial pair is sufficiently centered.

The key task, then, is to show how to update $(X^{(k)}, Z^{(k)})$ into $(X^{(k+1)}, Z^{(k+1)})$ such that (13) holds.

The algorithm depends on three parameters: $\nu \geq 1$ (which is used in the definition of the potential function), $\epsilon > 0$ is the value of the duality gap used to terminate the algorithm, and θ is a parameter satisfying $0 < \theta \leq 0.35$. (We shall see later that θ trades off numbers of iterations versus work per iteration.) Define $\delta = \theta - \log(1+\theta)$.

Primal-dual algorithm

given $W = X \oplus Z$ with $X \in \mathbf{Int}\,\mathcal{P} \cap (\mathcal{L} + D)$ and $Z \in \mathbf{Int}\,\mathcal{P} \cap (\mathcal{L}^\perp + C)$

repeat

1. *Find a suitable search direction.*
   Compute a $\delta W = \delta X \oplus \delta Z \in \mathcal{L} \oplus \mathcal{L}^\perp$ that satisfies
   $$\frac{\langle \delta W, \nabla\phi(W)\rangle}{\|W^{-1/2}\delta W W^{-1/2}\|} \geq \theta. \tag{15}$$
2. *Plane search.*
   Find p, $q \in \mathbf{R}$ such that $X - p\delta X \in \mathcal{P}$, $Z - q\delta Z \in \mathcal{P}$, and
   $$\phi\left(X - p\delta X, Z - q\delta Z)\right) \leq \phi(X, Z) - \delta. \tag{16}$$
3. *Update W.*
   Set $W = X \oplus Z := (X - p\delta X) \oplus (Z - q\delta Z)$.

until $\langle X, Z\rangle \leq \epsilon$.

We must clarify two points: first, how do we find a δW satisfying the condition (15), and second, how do we find p and q satisfying (16).

We first consider the search direction problem. It turns out that we can compute the direction $\delta W \in \mathcal{L} \oplus \mathcal{L}^\perp$ that maximizes the ratio on the left-hand side of (15) by solving a least squares problem. Define

$$\begin{aligned}
\delta W^N &= \operatorname*{argmin}_{\delta W \in \mathcal{L} \oplus \mathcal{L}^\perp} \left\| W^{1/2}\nabla\phi(W)W^{1/2} - W^{-1/2}\delta W W^{-1/2}\right\| \\
&= \operatorname*{argmin}_{\delta W \in \mathcal{L} \oplus \mathcal{L}^\perp} \left\| \rho W^{1/2}(C \oplus D)W^{1/2} - I - W^{-1/2}\delta W W^{-1/2}\right\|.
\end{aligned} \tag{17}$$

Then we have:

Theorem 3

$$\max_{\delta W \in \mathcal{L} \oplus \mathcal{L}^\perp} \frac{\langle \delta W, \nabla \phi(W) \rangle}{\|W^{-1/2} \delta W W^{-1/2}\|} = \frac{\langle \delta W^N, \nabla \phi(W) \rangle}{\|W^{-1/2} \delta W^N W^{-1/2}\|} = \left\| W^{-1/2} \delta W^N W^{-1/2} \right\| \geq \theta.$$

(This theorem is proved in the appendix.) Therefore we can always find a suitable search direction as required in step 1 of the algorithm by solving the least squares problem (17).

The direction δW^N can be interpreted as a Newton direction for a modified potential function. From (9) we see that ϕ is the sum of the convex function $F(X) + F(Z) - n \log n$ and the concave function $(n + \nu\sqrt{n}) \log\langle X, Z \rangle$. We first modify ϕ by linearizing the concave term at the point W. For strictly feasible $\tilde{W}$ we define

$$\varphi(\tilde{W}) = (n + \nu\sqrt{n}) \log\langle X, Z \rangle + \rho\langle \tilde{X} - X, Z \rangle + \rho\langle X, \tilde{Z} - Z \rangle + F(\tilde{X}) + F(\tilde{Z}) - n \log n$$

(X and Z, and therefore ρ, are fixed here). This modified potential function is convex. Moreover, since we have replaced a concave term by its linearization, we have $\varphi(W) = \phi(W)$ and $\varphi(\tilde{W}) \geq \phi(\tilde{W})$ for all feasible $\tilde{W}$. It follows that an update that reduces the modified potential φ will reduce the potential ϕ even more.

Now consider the Newton step for φ at the point W, *i.e.*, the minimizer of the quadratic model of φ at the point W:

$$\operatorname*{argmin}_{\delta W \in \mathcal{L} \oplus \mathcal{L}^\perp} \left(\varphi(W) + \langle \nabla\varphi(W), \delta W \rangle + \tfrac{1}{2} \langle \delta W, \nabla^2 \varphi(W) \delta W \rangle \right). \tag{18}$$

Since $\nabla\varphi(W) = \nabla\phi(W)$ and $\nabla^2\varphi(W) = \nabla^2 F(W)$, we have, using (5), $\nabla^2\varphi(W)\delta W = W^{-1}\delta W W^{-1}$. Therefore (18) becomes:

$$\begin{aligned} & \operatorname*{argmin}_{\delta W \in \mathcal{L} \oplus \mathcal{L}^\perp} \left(\langle \nabla\phi(W), \delta W \rangle + \tfrac{1}{2} \left\| W^{-1/2} \delta W W^{-1/2} \right\|^2 \right) \\ = & \operatorname*{argmin}_{\delta W \in \mathcal{L} \oplus \mathcal{L}^\perp} \tfrac{1}{2} \left\| W^{1/2} \nabla\phi(W) W^{1/2} + W^{-1/2} \delta W W^{-1/2} \right\|^2. \end{aligned}$$

Comparing this expression to (17) we see that this Newton step is precisely $-\delta W^N$.

We can also give a geometric interpretation of the search direction condition (15) in terms the angle between δW and the Newton direction δW^N in an appropriate metric. We can rewrite (17) as

$$\delta W^N = \operatorname*{argmin}_{\delta W \in \mathcal{L} \oplus \mathcal{L}^\perp} \left\| W^{-1/2} \left(W \nabla\phi(W) W - \delta W \right) W^{-1/2} \right\|$$

i.e., δW^N is the projection of $W \nabla\phi(W) W$ onto $\mathcal{L} \oplus \mathcal{L}^\perp$ in the norm $\|W^{-1/2}(\,\cdot\,)W^{-1/2}\|$. From this we see that

$$\begin{aligned} \langle \delta W, \nabla\phi(W) \rangle &= \langle W^{-1/2} \delta W W^{-1/2}, W^{1/2} \nabla\phi(W) W^{1/2} \rangle \\ &= \langle W^{-1/2} \delta W W^{-1/2}, W^{-1/2} \delta W^N W^{-1/2} \rangle. \end{aligned} \tag{19}$$

Therefore,

$$\frac{\langle \delta W, \nabla \phi(W) \rangle}{\|W^{-1/2}\delta W W^{-1/2}\|} = \frac{\langle W^{-1/2}\delta W W^{-1/2}, W^{-1/2}\delta W^N W^{-1/2}\rangle}{\|W^{-1/2}\delta W W^{-1/2}\|}$$

which is the norm of δW^N times the cosine of the angle between δW and δW^N, using the inner product $\langle W^{-1/2}(\,\cdot\,)W^{-1/2}, W^{-1/2}(\,\cdot\,)W^{-1/2}\rangle$. Evidently, the maximum value of this ratio is $\|W^{-1/2}\delta W^N W^{-1/2}\|$, obtained by choosing $\delta W = \delta W^N$. Therefore, in order to prove Theorem 3, we will need to show $\|W^{-1/2}\delta W^N W^{-1/2}\| \geq \theta$.

Now we turn to the question of plane search.

Theorem 4 *Suppose δW satisfies (15). Define*

$$p = q = \frac{\theta}{(1+\theta)\,\|W^{-1/2}\delta W W^{-1/2}\|}.$$

Then $X - p\delta X \in \mathcal{P}$, $Z - q\delta Z \in \mathcal{P}$, and

$$\phi\,(X - p\delta X, Z - q\delta Z)) \leq \phi(X, Z) - \delta.$$

(This theorem is proved in the appendix.) We note that this choice of p and q is only used in the complexity analysis. In practice, an approximate minimization of the potential function over p and q typically yields a much larger reduction in ϕ than δ (see section 4).

3 Search direction via conjugate gradients

In this section we show how to use a conjugate gradient algorithm to solve the least squares problem (17). Perhaps more importantly, the conjugate gradient algorithm we describe has the following properties:

– The iterates δW_k generated by the conjugate gradients algorithm are all feasible, *i.e.*, $\delta W_k \in \mathcal{L} \oplus \mathcal{L}^\perp$.

– δW_k converges to δW^N.

– The ratio

$$\frac{\langle \delta W_k, \nabla \phi(W) \rangle}{\|W^{-1/2}\delta W_k W^{-1/2}\|}$$

is known during the conjugate gradient process, and increases monotonically to its maximum value (which by Theorem 3 exceeds θ).

An immediate consequence is that we can stop the conjugate gradient process as soon condition (15) holds.

The least squares problem (17) requires us to project the matrix

$$B = \rho W^{1/2}(C \oplus D)W^{1/2} - I$$

onto the subspace $W^{-1/2}(\mathcal{L} \oplus \mathcal{L}^\perp)W^{-1/2}$. This subspace has dimension $n(n+1)/2$ so it can be represented as the range of a linear function $\mathcal{A}$ which maps $\mathbf{R}^{n(n+1)/2}$ into $\mathcal{S} \oplus \mathcal{S}$. Equivalently, we pick a basis for $W^{-1/2}(\mathcal{L} \oplus \mathcal{L}^\perp)W^{-1/2}$, with $\mathcal{A}y \in W^{-1/2}(\mathcal{L} \oplus \mathcal{L}^\perp)W^{-1/2}$ corresponding to coordinates $y \in \mathbf{R}^{n(n+1)/2}$.

Then $W^{-1/2}\delta W^N W^{-1/2} = \mathcal{A}\hat{y}$ where $\hat{y}$ is the solution of the least-squares problem

$$\hat{y} = \text{argmin } \|\mathcal{A}y - B\|. \tag{20}$$

Let $\mathcal{A}^*$ be the adjoint of $\mathcal{A}$, *i.e.*, the mapping from $\mathcal{S} \oplus \mathcal{S}$ into $\mathbf{R}^{n(n+1)/2}$ defined by $\langle H, \mathcal{A}y\rangle = (\mathcal{A}^*H)^T y$ for all $H \in \mathcal{S} \oplus \mathcal{S}$ and all $y \in \mathbf{R}^{n(n+1)/2}$. Then (20) can be solved via the normal equations

$$\mathcal{A}^*\mathcal{A}y = \mathcal{A}^*B. \tag{21}$$

The solution of this linear set of equations forms the main computational effort of the algorithm. It can be solved by a variety of direct or iterative methods, but as in implementations of interior-point methods for linear programming, the conjugate gradient method is a good choice [16, 1, 27, 25]. This method has the desirable property of only involving consecutive evaluations of the linear mappings $\mathcal{A}$ and $\mathcal{A}^*$ at given points. In particular, there is no need to form the matrix $\mathcal{A}^*\mathcal{A}$ which is of size $n(n+1)/2$.

The following is the outline of the conjugate gradient algorithm as it can for instance be found in [33]. Alternatively, one might prefer to use the LSQR-method [33]. LSQR is theoretically equivalent to the conjugate gradient method, but has better numerical properties.

The residual $r_k = \mathcal{A}^*B - \mathcal{A}^*\mathcal{A}y_k$ is calculated recursively. Criteria for termination will be discussed below.

Conjugate gradient algorithm

$k = 0$; $y_0 = 0$; $s_0 = B$, $r_0 = \mathcal{A}^*B$
until convergence
  $k = k + 1$
  if $k = 1$
    $p_1 = r_0$
  else
    $\beta_k = r_{k-1}^T r_{k-1} / r_{k-2}^T r_{k-2}$
    $p_k = r_{k-1} + \beta_k p_{k-1}$
  end
  $q_k = \mathcal{A}p_k$
  $\alpha_k = r_{k-1}^T r_{k-1} / q_k^T q_k$
  $y_k = y_{k-1} + \alpha_k p_k$
  $s_k = s_{k-1} - \alpha_k q_k$
  $r_k = \mathcal{A}^* s_k$
end
$\delta W = W^{1/2}\,(\mathcal{A}y_k)\,W^{1/2}$

This conjugate gradient algorithm finds the solution in $n(n+1)/2$ steps. We note that in practice a suitable preconditioner should be used. A good preconditioner can considerably speed up convergence, especially in the presence of roundoff error. The choice of a preconditioner depends on the specific structure of the problem, so we will postpone this topic to section 5.

In the rest of this section we discuss termination criteria for the conjugate gradient algorithm. We have:

Theorem 5 *Let $\delta W_k = W^{1/2}(\mathcal{A}y_k)W^{1/2}$ be the approximation of δW^N obtained after k iterations of the conjugate gradient algorithm. Then δW_k is a feasible direction with the property*

$$\frac{\langle \delta W_k, \nabla\phi(W)\rangle}{\|W^{-1/2}\delta W_k W^{-1/2}\|} = \left\|W^{-1/2}\delta W_k W^{-1/2}\right\| = \left(\sum_{i=1}^{k} \alpha_i^2 \; p_i^T \mathcal{A}^* \mathcal{A} p_i\right)^{1/2}. \tag{22}$$

Proof. The directions p_k are conjugate:

$$\langle \mathcal{A}p_j, \mathcal{A}p_k\rangle = 0, \quad j = 1, \ldots, k-1.$$

Since we have $y_k = \sum_{i=1}^k \alpha_i p_i$, it follows that

$$\langle W^{-1/2}\delta W_k W^{-1/2}, W^{-1/2}\delta W_k W^{-1/2}\rangle = \langle \mathcal{A}y_k, \mathcal{A}y_k\rangle = \sum_{i=1}^{k} \alpha_i^2 \; \langle \mathcal{A}p_i, \mathcal{A}p_i\rangle,$$

which proves the second equality. In order to prove the first one, let $\hat{y}$ be the solution of the normal equations (21) and observe that

$$\langle W^{-1/2}\delta W_k W^{-1/2}, W^{-1/2}\delta W^N W^{-1/2}\rangle = \langle \mathcal{A}y_k, \mathcal{A}\hat{y}\rangle = \langle \mathcal{A}y_k, \mathcal{A}y_k\rangle + y_k^T r_k.$$

The assertion now follows from (19) and from the orthogonality of the kth residual r_k and the directions p_j, $j = 1, \ldots, k$ (see [19]). □

This theorem has the interesting consequence that the ratio needed in (15) is readily known, and is easily computed since $\langle \mathcal{A}p_k, \mathcal{A}p_k\rangle$ has to be computed anyway. It suggests basing a stopping criterion for the conjugate gradient algorithm on the quantity (22). One possibility is to terminate as soon as condition (15) is satisfied. The parameter θ however is usually a very conservative lower bound for $\|W^{-1/2}\delta W^N W^{-1/2}\|$, and often a much larger reduction in potential function can be obtained by continuing the iteration after condition (15) is satisfied. Running the algorithm to completion on the other hand usually requires an excessive number of iterations, but one typically observes that the quantity (22) levels off long before the end of the algorithm. Preliminary experience therefore suggests terminating as soon as the relative increase in (22) becomes sufficiently small and condition (15) is satisfied.

4 Step lengths via plane search

Instead of the damped Newton step of Theorem 4 it is in general more efficient to make a plane search, *i.e.*, to look for scalars p and q that minimize $\phi(X - p\delta X, Z - q\delta Z)$. This can be done very efficiently if we first compute the eigenvalues $\mu_1, \ldots, \mu_n$ of $X^{-1/2}\delta X X^{-1/2}$ and the eigenvalues $\nu_1, \ldots, \nu_n$ of $Z^{-1/2}\delta Z Z^{-1/2}$. The potential function can then be written as

$$\phi(p,q) = (n + \nu\sqrt{n})\log(c_1 + c_2 p + c_3 q) + f(p,q) - n\log n$$

where

$$c_1 = \langle X, Z\rangle, \quad c_2 = -\langle Z, \delta X\rangle, \quad c_3 = -\langle X, \delta Z\rangle$$

and f is the restriction of the barrier term to the plane, *i.e.*,

$$\begin{aligned} f(p,q) &= F(X - p\delta X) + F(Z - q\delta Z) \\ &= -\sum_{i=1}^n \log(1 - p\mu_i) - \sum_{i=1}^n \log(1 - q\nu_i) - \log\det(XZ). \end{aligned}$$

The derivatives of these functions are straightforward to compute. We will need the first derivatives of $\phi(p,q)$

$$\frac{\partial \phi(p,q)}{\partial p} = \frac{c_2(n + \nu\sqrt{n})}{c_1 + pc_2 + qc_3} + \sum_{i=1}^n \frac{\mu_i}{1 - p\mu_i} \tag{23}$$

$$\frac{\partial \phi(p,q)}{\partial q} = \frac{c_3(n + \nu\sqrt{n})}{c_1 + pc_2 + qc_3} + \sum_{i=1}^n \frac{\nu_i}{1 - q\nu_i} \tag{24}$$

and the second derivatives of $f(p,q)$,

$$\frac{\partial^2 f(p,q)}{\partial p^2} = \sum_{i=1}^n \frac{\mu_i^2}{(1 - p\mu_i)^2}, \qquad \frac{\partial^2 f(p,q)}{\partial q^2} = \sum_{i=1}^n \frac{\nu_i^2}{(1 - q\nu_i)^2}. \tag{25}$$

Note that once we have computed the eigenvalues μ_i and ν_i, $i = 1, \ldots, n$, we can compute these derivatives in $O(n)$ operations.

In order to minimize ϕ we apply damped Newton steps to a linearized approximation of ϕ (as is done in [29, 7]). At each iteration the concave term of the potential function is linearized around the current (p,q) and one damped Newton step is applied to this modified potential. This involves the steps.

- Compute the derivatives (23), (24) and (25),
- Compute δp and δq from

$$\delta p = \frac{\partial \phi}{\partial p} \Big/ \frac{\partial^2 f}{\partial p^2}, \quad \delta q = \frac{\partial \phi}{\partial q} \Big/ \frac{\partial^2 f}{\partial q^2},$$

- Define λ as

$$\lambda = \left(\frac{\partial^2 f(p,q)}{\partial p^2}\delta p^2 + \frac{\partial^2 f(p,q)}{\partial q^2}\delta q^2 \right)^{1/2}$$

– Then the next iterates of p and q are

$$p := p - \frac{1}{1+\lambda}\delta p, \qquad q := q - \frac{1}{1+\lambda}\delta q.$$

It can be shown that with this choice of step length, $X - p\delta X$ and $Z - q\delta Z$ remain feasible and p, q converge to the values that minimize $\phi(p,q)$.

There is no need to calculate the minimum of $\phi(p,q)$ very accurately. Ben Tal and Nemirovsky [7] suggest taking a fixed number of steps. An alternative is to continue the iteration until λ becomes sufficiently small.

The main cost of this scheme is in the initial computation of the eigenvalues μ_i and ν_i. Once these are known, each step in the plane search can be carried out at a cost of $O(n)$ operations.

5 Algorithm for Lyapunov inequalities

We now return to the more specific problem (1). We will need to make a technical assumption: there is at least one index j, $1 \le j \le L$ for which the mapping

$$H \to B_j H A_j + A_j^T H B_j^T$$

is invertible. This is true if and only if the matrix pencil $B_j - \lambda A_j^T$ is regular and its spectrum σ is disjoint from its neqative, $\sigma \cap (-\sigma) = \emptyset$ (see [15]). The assumption is justified, because in most applications of (1) P also has to be positive semidefinite itself. One can ensure this by adding an inequality with $A_k = 0.5I$, $B_k = I$ and $D_k = 0$, which of course gives us an invertible mapping.

Problem (1) can be written in the conic form (2) as:

$$\min \sum_{k=1}^{L} \langle C_k, X_k \rangle, \quad X_k = A_k P B_k + B_k^T P A_k^T + D_k \ge 0, \quad k = 1, \dots, L,$$

which is of the form (2) with $X = X_1 \oplus \cdots \oplus X_L \in \bigoplus_{k=1}^{L} \mathcal{S}_m$, $n = mL$, and

$$\mathcal{L} = \left\{ \bigoplus_{k=1}^{L} A_k P B_k + B_k^T P A_k^T \;\middle|\; P \in \mathcal{S}_m \right\}, \quad D = \bigoplus_{k=1}^{L} D_k, \quad C = \bigoplus_{k=1}^{L} C_k.$$

In this conversion one is free to choose the matrices C_k as long as

$$\sum_{k=1}^{L} B_k C_k A_k + A_k^T C_k B_k^T = E. \tag{26}$$

This is an underdetermined set of equations. Its solution is in general not unique, but it can be verified that the choice of the C_k's has no effect on the algorithm. A general way to solve (26) is to apply the conjugate gradient algorithm, which would produce the minimum norm solution. An easier but less general way is to select an index j for which $B_j C_j A_j + A_j^T C_j B_j^T = E$ is solvable (we assumed above that such a j exists), compute this C_j (by the algorithm in [20, 15]), and take $C_k = 0$, $k \neq j$.

For this problem it is more natural to take $\bigoplus_{k=1}^{L} \mathcal{S}_m$ as the ambient space instead of $\mathcal{S}_{mL}$. The orthogonal complement of $\mathcal{L}$ then becomes

$$\mathcal{L}^\perp = \left\{ \; Z = \bigoplus_{k=1}^{L} Z_k \;\middle|\; \sum_{k=1}^{L} A_k^T Z_k B_k^T + B_k Z_k A_k = 0 \; \right\}$$

and we have the dual problem

$$\min \sum_{k=1}^{L} \langle D_k, Z_k \rangle, \quad \sum_{k=1}^{L} A_k^T Z_k B_k^T + B_k Z_k A_k = E, \quad Z_k \geq 0, \;\; k = 1, \ldots, L.$$

The total amount of work required to solve the problem (1) depends on two factors: the number of iterations, and the number of operations required in one iteration. We have seen that in the worst case the number of iterations grows with m and L as $O(\sqrt{mL})$, but in practice it appears to increase much more slowly. This behavior will be observed in the numerical experiments of the next section, and is consistent with the experience of other researchers with potential reduction methods (see [30, 22]). Throughout this section, then, we concentrate on the work per iteration.

The dominating part there is the solution of the least squares system (20), as discussed in section 3. In section 3 the projections on $\mathcal{L}$ and $\mathcal{L}^\perp$ were lumped together in a symmetrical way. This is not a good idea computationally, since usually $m \ll L$ and therefore the dimension of $\mathcal{L}^\perp$ ($\dim \mathcal{L}^\perp = (L-1)m(m+1)/2$) is much larger than the dimension of $\mathcal{L}$ ($\dim \mathcal{L} = m(m+1)/2$). A closer look however shows that the computational load of both projections is of the same order.

5.1 Direct methods

We can first clarify why we prefer to solve (20) by the conjugate gradient method instead of direct methods.

The least-squares problem (17) consists of two parts that can be solved independently: $\delta W^N = \delta X^N \oplus \delta Z^N$ with

$$\delta X^N = \underset{\delta X \in \mathcal{L}}{\operatorname{argmin}} \left\| \rho X^{1/2} C X^{1/2} - I - X^{-1/2} \delta X X^{-1/2} \right\| \tag{27}$$

$$\delta Z^N = \underset{\delta Z \in \mathcal{L}^\perp}{\operatorname{argmin}} \left\| \rho Z^{1/2} D Z^{1/2} - I - Z^{-1/2} \delta Z Z^{-1/2} \right\|. \tag{28}$$

The primal step (27) amounts to solving

$$\delta P = \underset{P \in \mathcal{S}_m}{\operatorname{argmin}} \sum_{k=1}^{L} \left\| X_k^{-1/2} \left(\rho X_k C_k X_k - X_k - A_k P B_k - B_k^T P^T A_k^T \right) X_k^{-1/2} \right\|^2$$

$$= \underset{P \in \mathcal{S}_m}{\operatorname{argmin}} \sum_{k=1}^{L} \left\| \rho X_k^{1/2} C_k X_k^{1/2} - I - \tilde{A}_k P \tilde{B}_k - \tilde{B}_k^T P \tilde{A}_k^T \right\|^2 \tag{29}$$

where the normalized data $\tilde{A}_k = X_k^{-1/2} A_k$ and $\tilde{B}_k = B_k X_k^{-1/2}$. From δP we can compute the Newton step $\delta X^N = \delta X_1^N \oplus \cdots \oplus \delta X_L^N$:

$$\delta X_k^N = A_k \delta P B_k + B_k^T \delta P A_k^T, \quad k = 1, \ldots, L.$$

A direct way to solve the least squares problem would consist in writing (29) as

$$\min_{P \in \mathcal{S}_m} \sum_{k=1}^{L} \left\| \operatorname{vec}\left(\rho X_k^{1/2} C_k X_k^{1/2} - I\right) - \left(\tilde{B}_k^T \otimes \tilde{A}_k + \tilde{A}_k \otimes \tilde{B}_k^T\right) \operatorname{vec}(P) \right\|^2, \tag{30}$$

where $A \otimes B$ denotes the Kronecker product, and $\operatorname{vec}(X)$ is the column vector obtained by appending the columns of the matrix X to one another. Problem (30) is an overdetermined set of equations with $O(m^2 L)$ equations and $O(m^2)$ unknowns and therefore requires $O(m^6 L)$ operations to solve by a direct method.

The dual step (28) is equivalent to

$$\delta Z^N = \delta Z_1^N \oplus \cdots \oplus \delta Z_L^N =$$
$$\operatorname{argmin} \left\{ \sum_{k=1}^{L} \left\| Z_k^{-1/2} \left(\rho Z_k D_k Z_k - Z_k - \delta Z_k\right) Z_k^{-1/2} \right\|^2 \;\middle|\; \sum_{k=1}^{L} A_k^T \delta Z_k B_k^T + B_k \delta Z_k A_k = 0 \right\}.$$

This is the projection of the matrix $\rho ZDZ - Z$ on $\mathcal{L}^\perp$, but in the metric associated with Z, *i.e.*, $\|Z^{-1/2}(\,\cdot\,)Z^{-1/2}\|$. Since the dimension of $\mathcal{L}^\perp$ is so much larger than the dimension of $\mathcal{L}$ it is advantageous to compute this projection by subtracting from $\rho ZDZ - Z$ its projection on $Z\mathcal{L}Z$ (this is the orthogonal complement of $\mathcal{L}^\perp$ in the metric induced by Z). The projection on $Z\mathcal{L}Z$ amounts to calculating

$$\delta P = \underset{P \in \mathcal{S}_m}{\operatorname{argmin}} \sum_{k=1}^{L} \left\| \rho Z_k^{1/2} D_k Z_k^{1/2} - I - \tilde{A}_k P \tilde{B}_k - \tilde{B}_k^T P \tilde{A}_k^T \right\|^2 \tag{31}$$

where $\tilde{A}_k = Z_k^{1/2} A_k$ and $\tilde{B}_k = B_k Z_k^{1/2}$. From δP one can find $\delta Z^N = \delta Z_1^N \oplus \cdots \oplus \delta Z_L^N$:

$$\delta Z_k^N = \rho Z_k D_k Z_k - Z_k - Z_k A_k \,\delta P B_k Z_k - Z_k B_k^T \,\delta P A_k^T Z_k, \quad k = 1, \ldots, L.$$

Using Kronecker products, problem (31) can again be converted to an overdetermined set of equations, analogous to (30). The computational cost is therefore $O(m^6 L)$ as well.

The corresponding figure if the least-squares system is solved by the conjugate gradient method is derived in the next section and will turn out to be $O(m^5 L)$.

5.2 Conjugate gradients

The direct method outlined above ignores two basic properties of the problem. It computes an exact solution, which is more than needed in the algorithm. Secondly, it is hard to take into account the Kronecker structure in the equations (30). With the conjugate gradient algorithm this will be much easier. We make use of the fact that solving a Sylvester equation $AXB + B^T XA^T = C$ where all matrices are of order m only requires $O(m^3)$ operations although the number of unknowns in the matrix X is of $O(m^2)$ (see [20, 15]).

The subspace $\mathcal{L}$ is of dimension $m(m+1)/2$, its orthogonal complement $\mathcal{L}^\perp$ of dimension $m(m+1)(L-1)/2$. In the conjugate gradient algorithm we can therefore use a representation

$$\mathcal{A} : \bigoplus_{k=0}^{L-1} \mathcal{S}_m \longrightarrow \bigoplus_{k=1}^{2L} \mathcal{S}_m,$$

mapping each set of L matrices $Y = Y_0 \oplus Y_1 \oplus \cdots \oplus Y_{L-1}$ on an element of $W^{-1/2}(\mathcal{L} \oplus \mathcal{L}^\perp)W^{-1/2}$. This is more convenient than using one single parameter vector. Defining $Y_p = Y_0$ and $Y_d = Y_1 \oplus \cdots \oplus Y_{L-1}$, we can write $\mathcal{A}$ as

$$\mathcal{A}Y = \mathcal{A}_p Y_p \oplus \mathcal{A}_d Y_d$$

where the images of

$$\mathcal{A}_p : \mathcal{S}_m \to \bigoplus_{k=1}^{L} \mathcal{S}_m \quad \text{and} \quad \mathcal{A}_d : \bigoplus_{k=1}^{L-1} \mathcal{S}_m \to \bigoplus_{k=1}^{L} \mathcal{S}_m$$

are $X^{-1/2}\mathcal{L}X^{-1/2}$ and $Z^{-1/2}\mathcal{L}^\perp Z^{-1/2}$, respectively.

From the description of the conjugate gradient algorithm it should be clear that no explicit representation of these mappings is needed. It is sufficient to have an efficient way of evaluating them at any given point.

Evaluation of $\mathcal{A}_p$ and $\mathcal{A}_p^*$. Let the scaled data matrices be $\tilde{A}_k = X_k^{-1/2} A_k$ and $\tilde{B}_k = B_k X_k^{-1/2}$, $k = 1, \ldots, L$. If $Y_0 \in \mathcal{S}_m$ then $\mathcal{A}_p Y_0 = \tilde{X}_1 \oplus \cdots \oplus \tilde{X}_L$ with

$$\tilde{X}_k = \tilde{A}_k Y_0 \tilde{B}_k + \tilde{B}_k^T Y_0 \tilde{A}_k^T, \quad k = 1, \ldots, L.$$

If $\tilde{X} = \tilde{X}_1 \oplus \cdots \oplus \tilde{X}_L \in \bigoplus_{k=1}^L \mathcal{S}_m$ then

$$\mathcal{A}_p^* \tilde{X} = \sum_{k=1}^{L} \tilde{B}_k \tilde{X}_k \tilde{A}_k + \tilde{A}_k^T \tilde{X}_k \tilde{B}_k^T.$$

Evaluation of $\mathcal{A}_d$ and $\mathcal{A}_d^*$. Let the scaled data matrices be $\tilde{A}_k = Z_k^{1/2} A_k$ and $\tilde{B}_k = B_k Z_k^{1/2}$, $k = 1, \ldots, L$. Pick j, $1 \le j \le L$ for which the mapping

$$H \to \tilde{B}_j H \tilde{A}_j + \tilde{A}_j^T H \tilde{B}_j^T$$

is invertible. Let $Y_d = Y_1 \oplus \cdots \oplus Y_{L-1}$, then the computation of

$$\tilde{Z} = \tilde{Z}_1 \oplus \cdots \oplus \tilde{Z}_L = \mathcal{A}_d Y_d$$

involves the following steps

1. Let $\tilde{C} = \sum_{k=1}^{j-1}(\tilde{B}_k Y_k \tilde{A}_k + \tilde{A}_k^T Y_k \tilde{B}_k^T) + \sum_{k=j+1}^{L}(\tilde{B}_k Y_{k-1} \tilde{A}_k + \tilde{A}_k^T Y_{k-1} \tilde{B}_k^T)$,
2. Compute $\tilde{Z}_j$ from $\tilde{B}_j \tilde{Z}_j \tilde{A}_j + \tilde{A}_j^T \tilde{Z}_j \tilde{B}_j^T = -\tilde{C}$.
3. Set $\tilde{Z}_k = Y_k$ for $k \neq j$.

If $\tilde{Z} = \tilde{Z}_1 \oplus \cdots \tilde{Z}_L \in \bigoplus_{k=1}^{L} \mathcal{S}_m$, then $Y_1 \oplus \cdots \oplus Y_{L-1} = \mathcal{A}_d^* \tilde{Z}$ can be computed as follows.

1. Solve $\tilde{A}_j \tilde{P} \tilde{B}_j + \tilde{B}_j^T \tilde{P} \tilde{A}_j^T = \tilde{Z}_j$ for $\tilde{P}$,
2. $Y_k = \tilde{Z}_k - \tilde{A}_k \tilde{P} \tilde{B}_k - \tilde{B}_k^T \tilde{P} \tilde{A}_k^T$, $k = 1, \ldots, j-1$ and $k = j+1, \ldots, L$.

Each of these four mappings can be evaluated in $O(m^3 L)$ operations. To estimate the number of iterations needed in the conjugate gradient method observe that $\mathcal{A}_d^* \mathcal{A}_d$ is the identity transformation plus a term with rank $O(m^2)$. Therefore the entire mapping $\mathcal{A}^* \mathcal{A}$ is a rank-$O(m^2)$ modification of the identity, which implies that the conjugate gradient algorithm converges in $O(m^2)$ steps (see [19]). The overall complexity to run the conjugate gradient algorithm to completion is therefore $O(m^5 L)$. In other words, by exploiting the Kronecker structure of the equations, the conjugate gradient method leads to a reduction of the $O(m^6 L)$ bound of the previous section by one order of magnitude. Moreover, several researchers have reported that with a good preconditioner, the conjugate gradient algorithm typically converges in much less than N iterations if N is the number of unknowns (*e.g.*, $O(\sqrt{N})$ iterations, see [25]). In the present case, it is therefore not unreasonable to expect that the number of iterations in the conjugate gradient algorithm can be reduced to $O(m)$, especially since there is no need to run it until completion. This would bring the overall complexity down to $O(m^4 L)$.

We conclude this section with some remarks concerning preconditioning. The mapping $\mathcal{A}$ is not the only possible choice. Finding a good preconditioner is equivalent to selecting a mapping $\mathcal{A}$ such that $\mathcal{A}^* \mathcal{A}$ is close to the identity matrix. A simple method is as follows.

Evaluation of $\mathcal{A}_p$ and $\mathcal{A}_p^*$. We can precede the evaluation of $\mathcal{A}_p$ by a linear transformation $\mathcal{M} : \mathcal{S}_m \rightarrow \mathcal{S}_m$ computed by solving an equation of the form

$$Y_0 = M(\mathcal{M} Y_0) N + N^T (\mathcal{M} Y_0) M^T.$$

A heuristic choice that seems to work well in practice is to take $M = \tilde{A}_j$ and $N = \tilde{B}_j$ where j is the index for which $\|X_k^{-1}\|$ is maximal. Then $(\mathcal{A}_p \mathcal{M})^* (\mathcal{A}_p \mathcal{M})$ can be expected to be closer to the identity mapping.

Evaluation of $\mathcal{A}_d$ and $\mathcal{A}_d^*$. In a similar way one can force $\mathcal{A}_d^*\mathcal{A}_d$ to be close to the identity transformation by selecting the index j in the evaluation of $\tilde{\mathcal{L}}^\perp$ to be the index for which $\|Z_k\|$ is maximal.

We have no doubt that there are more sophisticated methods of preconditioning for this problem. On the other hand, we will see that this method is effective enough to bring the average number of conjugate gradient iterations required per iteration to $O(m)$, which roughly speaking is the square root of the number of variables.

6 Some numerical experiments

Here we will give the results of two numerical experiments that confirm the complexity estimate of $O(m^4L^{1.5})$ for the problem (1).

We consider two families of problems. Both are of the form

$$\min \mathbf{Tr}\, EP, \quad X_k = A_kP + PA_k^T + D_k \geq 0, \quad k = 1, \dots, L, \tag{32}$$

with dual problem

$$\min \sum_{k=1}^{L} \mathbf{Tr}\, D_kZ_k, \quad \sum_{k=1}^{L} Z_kA_k + A_k^TZ_k = E, \quad Z_k \geq 0, \;\; k = 1, \dots, L. \tag{33}$$

In the first family, the problem data are randomly generated. The second family of problems is derived from an application from control theory.

6.1 A family of random problems

Random data were generated as follows. The Lth inequality is $2P \geq I$, or in other words $A_L = I$, $D_L = -I$. To get the other inequalities we first form a random positive definite matrix $A_k^{\mathrm{sym}} = V_k^T\Lambda_kV_k$. Here Λ_k is a diagonal matrix with diagonal elements uniformly distributed in the interval $(0, 1)$. The matrix V_k is an orthonormal matrix drawn from a uniform distribution on the $m \times m$ orthogonal matrices. To A_k^{sym} we add a random skew-symmetric matrix $A_k^{\mathrm{ss}} = S_k - S_k^T$ where S_k has elements normally distributed with zero mean and standard deviation one. The matrix A_k is then obtained as $A_k = A_k^{\mathrm{sym}} + A_k^{\mathrm{ss}}$. Furthermore we take $D_k = 0$, $k = 1, \dots, L-1$ and $E = \sum_{i=1}^{L}(A_i + A_i^T)$. With this choice one has obvious primal and dual feasible solutions

$$P = I, \quad X_k = A_k + A_k^T, \quad Z_k = I, \;\; k = 1, \dots, L.$$

The parameter θ was chosen to be 0.35, but in fact the conjugate gradient iterations were continued until in addition the ratio (15) was no longer increasing. The parameter ν was set to $\nu = 100$, although for this problem family, the value of ν did not greatly affect the results (see section 6.3). The stopping criterion was a reduction of the initial duality gap by a factor

m	L	Total no. of c.g. iterations	No. of outer iterations	Average no. of c.g. iterations
5	10	25	6.3	4.4
5	30	33	7.0	4.7
5	50	33	7.0	4.6
5	70	35	7.1	4.9
10	10	74	11.7	8.0
10	30	67	7.1	9.4
10	50	79	7.5	10.6
10	70	79	7.7	10.3
15	10	92	7.0	13.2
15	30	114	7.4	15.3
15	50	153	8.0	19.1
15	70	142	8.0	17.7
20	10	124	7.0	17.7
20	30	153	8.3	18.6
20	50	208	8.0	26.0
20	70	159	8.2	19.5

Table 1: Results of the first experiment: total number of conjugate gradient iterations, number of outer iterations of the primal-dual algorithm, and average number of conjugate gradient iterations per step of the primal-dual algorithm. Each figure is the average of ten randomly generated problems.

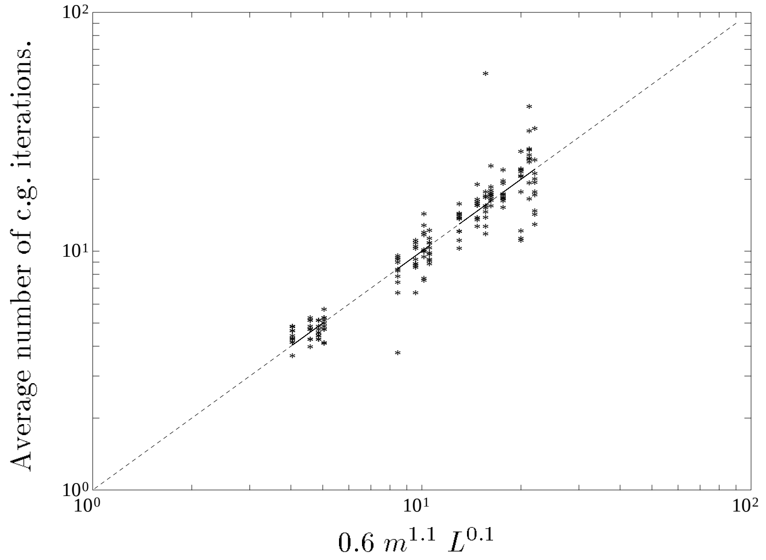

Figure 1: First experiment. Average number of conjugate gradient iterations versus $0.6\ m^{1.1}\ L^{0.1}$.

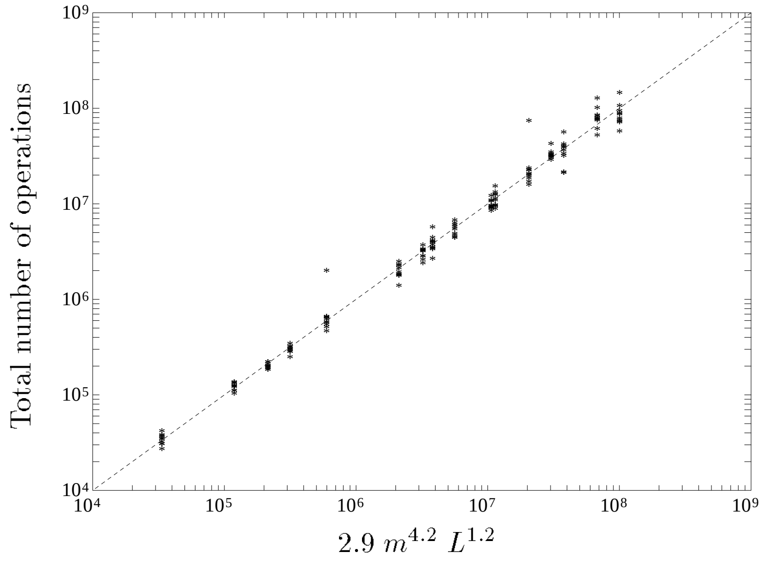

Figure 2: First experiment. Total number of operations (up to a constant factor) vs. $2.9\ m^{4.2}\ L^{1.2}$.

of 1000. The results are listed in Table 1. Each entry is an average over ten instances of the problem. For the largest problem the number of unknowns is $m(m+1)/2 = 210$ and in each iteration the least-squares problem has size 14700×210.

From table 1 several interesting conclusions can be drawn.

1. The number of outer iterations in the primal-dual method is almost constant or at least growing very slowly. Similar observations have been made for other potential reduction methods (see *e.g.*, [30, 22]). This implies that the determining factor in the overall complexity is the computational cost of one iteration.

2. The number of conjugate gradient iterations grows linearly with m, which is the square-root of the number of variables in the least squares system that is solved. This is consistent with results reported in [25].

   A least squares fit of these points to a curve of the form $\alpha n^{\beta} L^{\gamma}$ results in $\alpha = 0.56$, $\beta = 1.06$ and $\gamma = 0.11$. Figure 1 shows a scatter plot of the average number of conjugate gradient iterations per outer iteration versus $0.6m^{1.1}L^{0.1}$ over the 160 experiments.

3. Recall that the number of operations required per conjugate gradient iteration is $O(m^3L)$. The total number of operations is therefore proportional to m^3L times the total number of conjugate gradient iterations. (Each conjugate gradient iteration requires $4L$ $m \times m$ matrix multiplies and the solution of four $m \times m$ Sylvester equations, so in fact the constant here is not too big, provided L is not very small.)

   If we fit this number to a function of the form $\alpha m^{\beta} L^{\gamma}$ we obtain $\alpha = 2.9$, $\beta = 4.2$ and $\gamma = 1.2$. Figure 2 shows a scatter plot of m^3L times the total number of conjugate gradient iterations required for convergence versus $2.9m^{4.2}L^{1.2}$ over the 160 experiments.

   Note that $2.9m^{4.2}L^{1.2}$ predicts the effort typically within $\pm 25\%$. Even in the worst case, the error is only a factor of 2 or 3.

   It is also interesting to compare this with the worst-case complexity of the ellipsoid method when applied to the same problem. The number of iterations in the ellipsoid method is $O(m^4)$ (*i.e.*, the number of unknows squared, see [23]), and in this problem the cost of each iteration is $O(m^3L)$. This results in an overall complexity of $O(m^7L)$, much higher than in the interior-point method described here.

6.2 A family of problems from system theory

The same experiments are now repeated for a family of problems from system theory. We first give some general background for the problems. We consider a time varying linear system described by

$$\frac{dx}{dt} = A(t)x(t), \quad y = c^T x \tag{34}$$

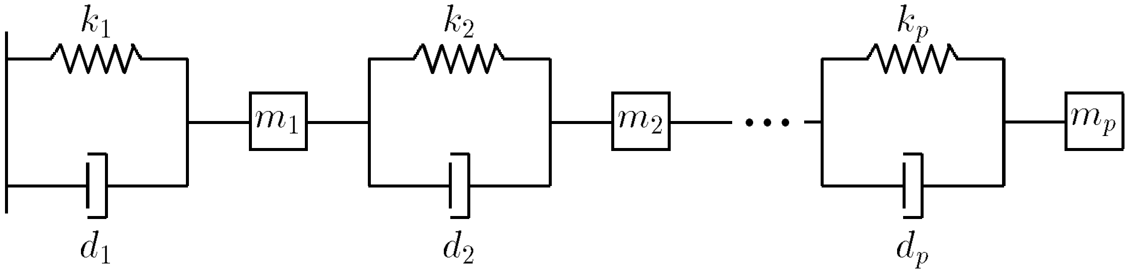

Figure 3: Mechanical system with p masses, springs and dampers. The masses and dampers are constant with unit value. Some of the springs can vary arbitrarily over the range $[0.9, 1.1]$; the remaining springs are fixed with unit value.

where the $m \times m$ matrix $A(t)$ is known to lie in the convex hull of $L-1$ matrices:

$$A(t) \in \mathbf{Co}\{A_1, \ldots, A_{L-1}\} \tag{35}$$

for all $t \in \mathbf{R}$. (This is called a linear differential inclusion.)

We want to bound the maximum value of the output y given a bound on the norm of the initial state, using a quadratic Lyapunov function. The following theorem is readily derived from the results of [9]:

Theorem 6 *Suppose that (34) and (35) hold, and in addition $\|x(0)\| \leq 1$. Then for all $t \geq 0$, $|y(t)| \leq \alpha$ where*

$$\alpha = \min_{\substack{A_iP + PA_i^T \leq 0, \; i=1,\ldots,L-1 \\ P \geq I}} \mathbf{Tr}\, cc^T P.$$

(In fact, α is the best bound that can be obtained using a quadratic Lyapunov function.) Therefore, the problem of computing α is of the form (32).

Now we can describe the specific family of problems we consider. Figure 3 shows a mechanical system consisting of p masses, springs, and dampers. The masses and dampers are constant with value one. Some of the springs, however, can vary arbitrarily with time, between the limits 0.9 and 1.1. The remaining springs are fixed with value one. The output y that concerns us is the displacement of the pth mass. For this problem, we have $m = 2p$ and $L = 2^r + 1$, where r is the number of springs that vary.

In our experiments we compute the bound α for several values of p and r (which results in several values of m and L). For these problems there are no obvious feasible solutions that can be used to start the algorithm. In order to use initial points that could be fairly compared across different values of m and L, we take the primal and dual central pair with duality gap equal to mL as the starting points, *i.e.*, $X^{(0)} = X^*(mL)$ and $Z^{(0)} = Z^*(mL)$. The stopping criterion is again a reduction of the initial duality gap by a factor of 1000 (*i.e.*, $\epsilon = 0.001mL$).

$m = 2p$	L	Indices of the time-varying springs	Total no. of c.g. iterations	No. of outer iterations	Average no. of c.g. iterations
4	3	1	79	9	8.8
4	5	1,2	93	10	9.3
6	3	1	163	11	14.8
6	5	1,2	224	12	18.7
6	9	1,2,3	230	12	19.2
8	3	1	200	12	16.7
8	5	1,2	259	12	21.6
8	9	1,2,3	377	13	29.0
8	17	1,2,3,4	641	16	40.1
10	3	1	213	13	16.4
10	5	1,2	320	14	22.9
10	9	1,2,3	489	14	34.9
10	17	1,2,3,4	1040	18	57.8
12	3	1	251	14	17.9
12	5	1,2	273	14	19.5
12	9	1,2,3	381	16	23.8
12	17	1,2,3,4	858	18	47.7

Table 2: Results of the second experiment: total number of conjugate gradient iterations, number of outer iterations of the primal-dual algorithm, and average number of conjugate gradient iteriations per step of the primal-dual algorithm.

For this family of problems, the parameter ν has a greater effect on convergence than for the randomly generated problems. Large values of ν (*e.g.*, $\nu = 100$) resulted in slower convergence than smaller values. Table 2 give the results for the case $\nu = 2$. For this family of problems, the best fit to the total number of operations is $9.8m^{4.0}L^{1.7}$. Figure 4 shows that this expression predicts the total number of operations very well: the average prediction error is about 20%.

6.3 Influence of the parameter ν

In this section we study the influence of the parameter ν on the primal dual algorithm. We take an instance of the random problem of section 6.1 with $m = 4$ and $L = 2$, while varying the parameter ν and the initial pair. One of the initial points lies on the central path, while the other four lie off the central path at increasing distances.

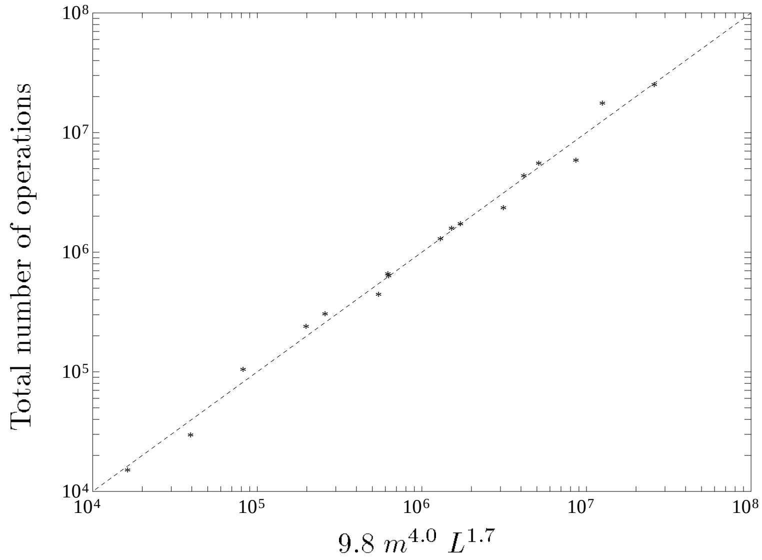

Figure 4: Second experiment. Total number of operations (up to a constant factor) vs. $9.8\ m^{4.0}\ L^{1.7}$.

Figure 5 shows the trajectories of the primal dual algorithm in a two-dimensional plane showing duality gap $\langle X, Z\rangle$ and the function $\psi(X, Z)$. Thus the horizontal axis shows the duality gap on a logarithmic scale, and the vertical axis shows the deviation from centrality (measured in units of 3.26 Newton steps, from our interpretation (7)). From (8), the level curves of ϕ are straight lines in these plots, with a slope that depends on ν. These are shown with the dotted lines.

Some observations from these plots are:

- After the first iteration, the potential function decreases quite linearly. The typical reduction in ϕ per iteration increases with ν.

- Increasing ν places more weight on cost reduction (versus centering) and will at first speed up convergence. (For large values of ν, the algorithm behaves like Dikin's affine scaling method [11]; convergence slows down again, because the iterates come too close to the boundary.)

- The paths approach a limiting value of ψ that depends on ν. In other words, the centrality of the iterates eventually remains nearly constant: for $\nu = 1$, the iterates eventually remain in or very near the region of quadratic convergence surrounding the central path; for $\nu = 100$, the iterates remain about 33 Newton steps away from the region of quadratic convergence.

Once this limiting value of ψ has been reached, the decrease in ϕ at each iteration is entirely due to decrease in duality gap. We do not entirely understand the mechanism involved here, but it seems to be linked to the use of an almost exact plane search.

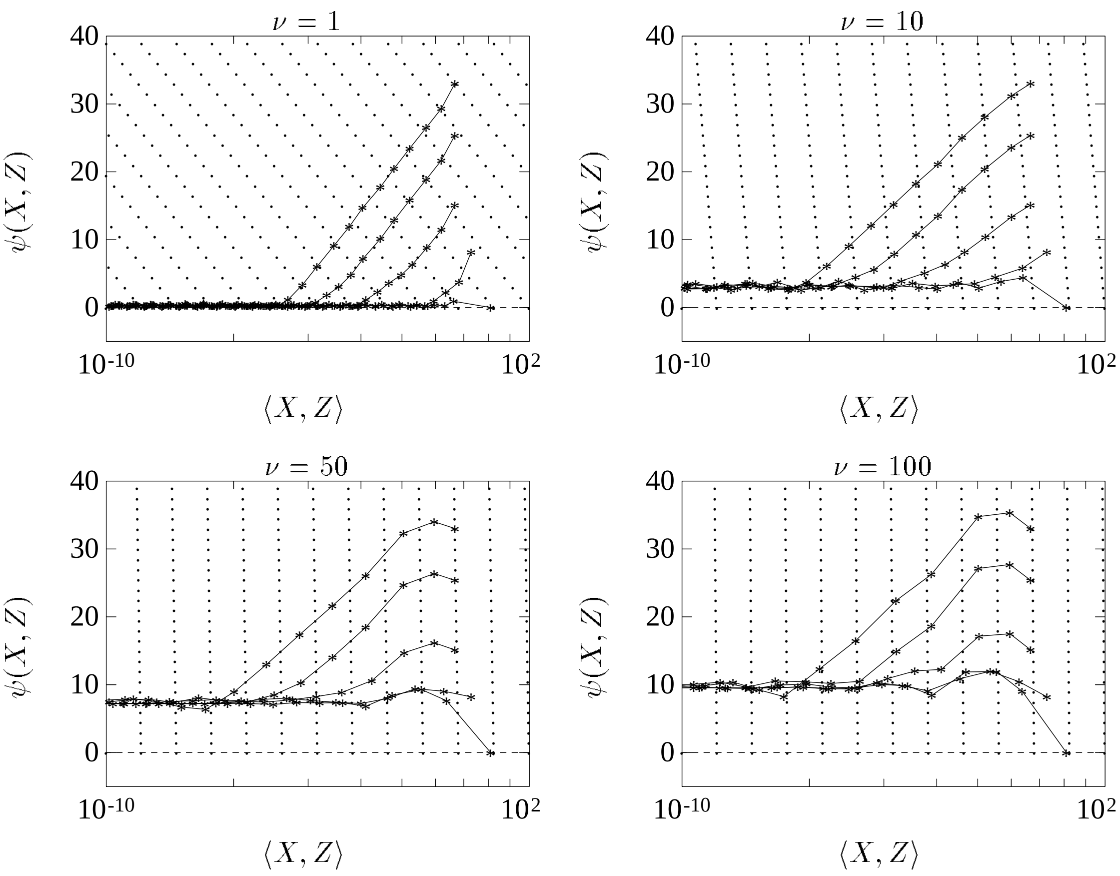

Figure 5: Convergence of the algorithm for four values of ν. Each plot shows the paths followed by the algorithm for five different initial points (full lines). The dashed line corresponds to the central path, $\psi(X, Z)$ measures the deviation from the central path and the dotted lines are the level curves of the potential function $\phi(X, Z)$.

7 Conclusions

7.1 Extension to Riccati inequalities

The ideas presented here can be extended to matrix inequalities other than those of the Lyapunov type. As an example consider a problem constrained by L algebraic Riccati inequalities:

$$\min \mathbf{Tr}\, EP, \quad -A_k^T P - PA_k - PB_k R_k^{-1} B_k^T P - Q_k \geq 0, \quad k = 1, \ldots, L,$$

where E, A_k, $Q_k \in \mathbf{R}^{m\times m}$, $B_k \in \mathbf{R}^{m\times l}$, $R_k \in \mathbf{R}^{l\times l}$, $k = 1, \ldots, L$ and Q_k and R_k are symmetric with R_k positive definite. As mentioned in section 1, we can express this problem as

$$\min \mathbf{Tr}\, EP, \quad \text{s.t.} \quad \bar{A}P\bar{B}_k + \bar{B}_k^T P \bar{A}^T + \bar{D}_k \geq 0, \;\; k = 1, \ldots, L. \tag{36}$$

where

$$\begin{aligned} \bar{A} &= \begin{bmatrix} I \\ 0 \end{bmatrix}, \\ \bar{B}_k &= \begin{bmatrix} -A_k & B_k \end{bmatrix}, \\ \bar{D}_k &= \begin{bmatrix} -Q_k & 0 \\ 0 & R_k \end{bmatrix}. \end{aligned}$$

For this problem the forward mapping $\mathcal{A}_p$ is defined as

$$\mathcal{A}_p \tilde{P} = \bigoplus_{k=1}^{L} \tilde{A}_k \tilde{P} \tilde{B}_k + \tilde{B}_k^T \tilde{P} \tilde{A}_k^T$$

where $\tilde{A}_k = X_k^{-1/2} \bar{A}$, $\tilde{B}_k = \bar{B}_k X_k^{-1/2}$, $k = 1 \ldots, L$. The adjoint $\mathcal{A}_p^*$ is

$$\mathcal{A}_p^* \bigoplus_{k=1}^{L} \tilde{Z}_k = \tilde{A}_k^T \tilde{Z}_k \tilde{B}_k^T + \tilde{B}_k \tilde{Z}_k \tilde{A}_k.$$

Problem (36) can now be converted into the conic form by taking as subspaces $\mathcal{L}$ and $\mathcal{L}^\perp$,

$$\mathcal{L} = \left\{ \mathcal{A}_p \tilde{P} \,\middle|\, \tilde{P} \in S_m \right\}, \quad \mathcal{L}^\perp = \left\{ \tilde{Z} \in \bigoplus_{k=1}^{L} \mathcal{S}_{m+l} \,\middle|\, \mathcal{A}_p^* \tilde{Z} = 0 \right\}.$$

One can choose for D

$$D = \bigoplus_{k=1}^{L} \begin{bmatrix} -Q_k & 0 \\ 0 & R_k \end{bmatrix},$$

and for C any matrix of the form $C \in \bigoplus_{k=1}^{L} \mathcal{S}_{m+l}$ for which $\mathcal{A}_p^* C = E$.

We may assume that $l \leq m$. Then we see that both $\mathcal{A}_p$ and $\mathcal{A}_p^*$ can be computed at given points in $O(m^3 L)$ operations. An explicit representation $\mathcal{A}_d$ for the subspace $\mathcal{L}^\perp$ can be derived in similar way as we did for Lyapunov inequalities. Both $\mathcal{A}_d$ and its adjoint can be evaluated in $O(m^3 L)$ operations. This implies that one can solve optimization problems over Riccati inequalities at the same low cost as problems with Lyapunov inequalities.

Efficient algorithms for problem (36) allow us to extend traditional, single-model LQR (Linear Quadratic Regulator) controllers to the multiple model case. This will be the subject of a forthcoming paper.

We note that solving the optimization problem (36) requires the solution of Lyapunov equations, but does not require the solution of any Riccati equations. In fact, by taking $E = I$ and $L = 1$, the solution $P^{\rm opt}$ of the optimization problem (36) is the solution of the algebraic Riccati equation $A^TP + PA + PBR^{-1}B^TP + Q = 0$ for which $A + BR^{-1}B^TP$ is stable,*i.e.*, all eigenvalues have negative real part. The primal dual algorithm requires effort $O(m^4)$ to compute $P^{\rm opt}$, which is only one order higher than conventional algorithms (see, *e.g.*, [5]). (See [9] for a complete discussion of the primal-dual matrix inequality formulation of the classical LQR problem.)

7.2 Possible modifications of the method

The interior-point method itself can also be extended in several directions. One possibility is to combine the algorithm with Ye's method [36] and its dual. The directions used in Ye's method can be computed from the Newton directions used in the present paper, and the two-dimensional plane search would then be replaced by a four-dimensional search. This algorithm coincides with the extension of Ye's method suggested by Nesterov and Nemirovsky [30, Ch.4].

We have also already noted the important role the parameter ν can play. The optimal value of ν is clearly problem dependent, so strategies for choosing it or adapting it during the algorithm are certainly worth investigating (see, for example, [34]).

7.3 Extension to more general problems

In this section we make some brief remarks about how the algorithm presented in this paper can be extended to the more general conic formulation of Nesterov and Nemirovsky [30]. In this general context the cones in the primal and dual problems (2) and (3) are different: the cone in the dual problem is the dual of the cone in the primal problem. (We have not made this distinction since $\mathcal{P}$ is self-dual.)

The primal-dual algorithm applies to cones admitting a θ-logarithmically homogeneous self-concordant barrier (see [30]). The parameter θ of the barrier is equal to n in our case, and determines the worst-case complexity (which is $O(\sqrt{\theta})$ in the general case).

If F and F^* are the barriers for the cone and its dual, the primal-dual potential function is defined as

$$(\theta + \nu\sqrt{\theta}) \log\langle X, Z\rangle + F(X) + F^*(Z).$$

Again the term $\theta \log\langle X, Z\rangle + F(X) + F^*(Z)$ is bounded below, and attains its minimum for points X, Z with $Z = -(\alpha/\theta)\nabla F(X)$, $X = -(\alpha/\theta)\nabla F(Z)$, which characterizes the central path.

The complexity analysis of the algorithm given here remains valid in the general case. The generalizations of the relevant results either follow immediately from the general theory in [30] or can readily be extended from the results given in this paper.

8 Acknowledgments

We are indebted to Arkadii Nemirovskii, Florian Jarre, Michael Saunders, Michael Overton, Jean-Pierre Haeberly, Gene Golub and Bart De Moor for useful discussions and comments on an earlier version of this paper.

References

[1] I. Adler, M. G. C. Resende, G. Veiga, and N. Karmarkar, "An implementation of Karmarkar's algorithm for linear programming", Mathematical Programming 44 (1989) 297–335.

[2] F. Alizadeh, "Combinatorial optimization with interior point methods and semi-definite matrices", PhD thesis, University of Minnesota (October 1991).

[3] F. Alizadeh, "Optimization over the positive-definite cone: interior point methods and combinatorial applications", in: P. Pardalos, ed., Advances in Optimization and Parallel Computing (North-Holland, 1992).

[4] K. M. Anstreicher, "A combined phase I – phase II scaled potential algorithm for linear programming", Mathematical Programming 52 (1991) 429–439.

[5] W. F. Arnold and A. J. Laub, "Generalized eigenproblem algorithms and software for algebraic Riccati equations", Proceedings of the IEEE 72 (1984) 1746–1754.

[6] R. Bellman and K. Fan, "On systems of linear inequalities in Hermitian matrix variables", in: V. L. Klee, ed., Convexity, Volume 7 of The Proceedings of Symposia in Pure Mathematics, (American Mathematical Society, 1963) pp. 1–11.

[7] A. Ben Tal and A. Nemirovskii, "Interior point polynomial time method for truss topology design", Technical Report 3/92, Faculty of Industrial Engineering and Management, Technion, Israel Institute of Technology (1992).

[8] S. Boyd and L. El Ghaoui, "Method of centers for minimizing generalized eigenvalues", Linear Algebra and Applications (special issue on Numerical Linear Algebra Methods in Control, Signals and Systems) 188 (1993) 63–111.

[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, 1994).

[10] D. Den Hertog, C. Roos, and T. Terlaky, "A large-step analytic center method for a class of smooth convex programming problems", SIAM Journal on Optimization 2 (1992) 55–70.

[11] I. Dikin, "Iterative solution of problems of linear and quadratic programming", Soviet Mathematics Doklady 8 (1967) 674–675.

[12] R. Fletcher, "A nonlinear programming problem in statistics (educational testing)", SIAM Journal on Scientific and Statistical Computing 2 (1981) 257–267.

[13] R. Fletcher, "Semidefinite matrix constraints in optimization", SIAM Journal on Control and Optimization 23 (1985) 493–513.

[14] R. M. Freund, "A potential-function reduction algorithm for solving a linear program directly from an infeasible warm start", Mathematical Programming 52 (1991) 441–466.

[15] J. D. Gardiner, A. J. Laub, J. J. Amato, and C. B. Moler, "Solution of the Sylvester matrix equation $AXB^T + CXD^T = E$", ACM Transactions on Mathematical Software 18 (1992) 223–231.

[16] P. E. Gill, W. Murray, M. A. Saunders, J. A. Tomlin, and M. H. Wright, "On projected Newton barrier methods for linear programming and an equivalence to Karmarkar's projective method", Mathematical Programming Studies 36 (1986) 183–209.

[17] D. Goldfarb and S. Mehrotra, "A relaxed version of Karmarkar's method", Mathematical Programming 40 (1988) 289–315.

[18] D. Goldfarb and S. Mehrotra, "A self-correcting version of Karmarkar's algorithm", SIAM Journal on Numerical Analysis 26 (1989) 1006–1015.

[19] G. Golub and C. Van Loan, Matrix Computations (Johns Hopkins University Press, Baltimore, 1989).

[20] G. Golub, S. Nash, and C. Van Loan, "A Hessenberg-Schur method for the matrix problem $AX + XB = C$", IEEE Transactions on Automatic Control 24 (1979) 909–913.

[21] C. C. Gonzaga, "Path-following methods for linear programming", SIAM Review 34 (1992) 167–224.

[22] C. C. Gonzaga and M. J. Todd, "An $O(\sqrt{n}L)$-iteration large-step primal-dual affine algorithm for linear programming", SIAM Journal on Optimization 2 (1992) 349–359.

[23] M. Grötschel, L. Lovász, and A. Schrijver, Geometric Algorithms and Combinatorial Optimization, Volume 2 of *Algorithms and Combinatorics* (Springer-Verlag, 1988).

[24] F. Jarre, "An interior-point method for minimizing the maximum eigenvalue of a linear combination of matrices", SIAM Journal on Control and Optimization 31 (1993) 1360–1377

[25] K. Kim and J. L. Nazareth, "Implementation of a primal null-space affine scaling method and its extensions", Technical report, Department of Pure and Applied Mathematics, Washington State University (1992).

[26] I. J. Lustig, "Feasibility issues in a primal-dual interior-point method for linear programming", Mathematical Programming 49 (1991) 145–162.

[27] S. Mehrotra, "Implementations of affine scaling methods: approximate solutions of systems of linear equations using preconditioned conjugate gradient methods, ORSA Journal on Computing 4 (1992) 103–118.

[28] Renato D. C. Monteiro and Ilan Adler, "Interior path-following primal-dual algorithms. part II: convex quadratic programming", Mathematical Programming 44 (1989) 43–66.

[29] Yu. Nesterov and A. Nemirovsky, Optimization over positive semidefinite matrices: Mathematical background and user's manual, (USSR Acad. Sci. Centr. Econ. & Math. Inst., 32 Krasikova St., Moscow 117418 USSR, 1990).

[30] Yu. Nesterov and A. Nemirovsky, Interior-point polynomial methods in convex programming, Volume 13 of *Studies in Applied Mathematics* (SIAM, Philadelphia, PA, 1994).

[31] M. Overton, "On minimizing the maximum eigenvalue of a symmetric matrix", SIAM Journal on Matrix Analysis and Applications 9 (1988) 256–268.

[32] M. Overton, "Large-scale optimization of eigenvalues", SIAM Journal on Optimization 2 (1992) 88–120.

[33] C. C. Paige and M. S. Saunders, "LSQR: An algorithm for sparse linear equations and sparse least squares", ACM Transactions on Mathematical Software, 8 (1982) 43–71.

[34] J.-P. Vial, "Computational experience with a primal-dual interior-point method for smooth convex programming", Technical report, Département d'Economie Commerciale et Industrielle, Université de Genève (1992).

[35] G. A. Watson, "Algorithms for minimum trace factor analysis", SIAM Journal on Matrix Analysis and Applications 13 (1992) 1039–1053.

[36] Y. Ye, "An $O(n^3L)$ potential reduction algorithm for linear programming", Mathematical Programming 50 (1991) 239–258.

A Proofs

A.1 Proof of Theorem 1

The feasible sets in problem (6) are bounded and convex with nonempty relative interior. Moreover the objective function $F(X) + F(Z)$ is a strictly convex function, and therefore the central points are the unique points that satisfy the optimality conditions

$$-X^{-1} + \mu C \in \mathcal{L}^{\perp}, \quad -Z^{-1} + \mu D \in \mathcal{L},$$

or,

$$(1/\mu)X^{-1} \in \mathcal{L}^{\perp} + C, \quad (1/\mu)Z^{-1} \in \mathcal{L} + D,$$

for some value of the Lagrange multiplier μ, along with the feasibility conditions

$$X \in \mathcal{P} \cap (\mathcal{L} + D), \quad Z \in \mathcal{P} \cap (\mathcal{L}^{\perp} + C), \quad \langle X, Z \rangle = \alpha.$$

Now note that if X, Z satisfy the above conditions, then so does the pair $(1/\mu)Z^{-1}, (1/\mu)X^{-1}$. From uniqueness we conclude that

$$X^*(\alpha) = (1/\mu)Z^*(\alpha)^{-1}, \quad Z^*(\alpha) = (1/\mu)X^*(\alpha)^{-1},$$

and therefore $\mu = n/\alpha$.

$\square$

A.2 Proof of Theorem 3

Lemma 1 *If $x \in \mathbf{R}^n$ and $\|x\| < 1$, then $\sum_{i=1}^n x_i - \sum_{i=1}^n \log(1 + x_i) \le -\|x\| - \log(1 - \|x\|)$.*

Proof. (see [21])

$$\begin{aligned}
\sum_{i=1}^n x_i - \sum_{i=1}^n \log(1 + x_i) &= \sum_{i=1}^n x_i - \sum_{i=1}^n x_i + \sum_{k=2}^{\infty} \frac{(-1)^k}{k} \sum_{i=1}^n x_i^k \\
&\le \sum_{k=2}^{\infty} \frac{1}{k} \sum_{i=1}^n (x_i^2)^{k/2} \\
&\le \sum_{k=2}^{\infty} \frac{1}{k} \|x\|^k = -\log(1 - \|x\|) - \|x\|,
\end{aligned}$$

where the second inequality follows from the fact that $\sum_{i=1}^n a_i^p \le (\sum_{i=1}^n a_i)^p$ if $a_i \ge 0$, $i = 1, \ldots, n$ and $p \ge 1$. $\square$

Lemma 2 *If $\alpha > 0$ and X and Z are strictly feasible, then*

$$(n/\alpha)\langle X, Z \rangle + F(X) + F(Z) \ge n + n \log(n/\alpha).$$

If in addition $\|(n/\alpha)X^{1/2}ZX^{1/2} - I\| = \epsilon < 1$, then

$$(n/\alpha)\langle X, Z \rangle + F(X) + F(Z) \le n + n \log(n/\alpha) - \epsilon - \log(1 - \epsilon) \tag{37}$$

Proof. Let $\mu_1, \dots, \mu_n$ be the eigenvalues of $(n/\alpha)XZ - I$. Strict feasibility of X and Z implies $\mu_i > -1$, $i = 1, \dots, n$. We have

$$\begin{aligned}(n/\alpha)\langle X, Z\rangle + F(X) + F(Z) &= n + n\log(n/\alpha) + \sum_{i=1}^n \mu_i - \sum_{i=1}^n \log(1+\mu_i)\\ &\geq n + n\log(n/\alpha),\end{aligned}$$

because $\mu_i - \log(1+\mu_i) \geq 0$ if $\mu_i > -1$. Inequality (37) is an immediate consequence of Lemma 1 if $\|(n/\alpha)X^{1/2}ZX^{1/2} - I\| = (\sum_{i=1}^n \mu_i^2)^{1/2} = \epsilon < 1$. □

Proof of Theorem 3. We will prove the theorem by showing that

$$\max\left\{ \left\|X^{-1/2}\delta X^N X^{-1/2}\right\|, \left\|Z^{-1/2}\delta Z^N Z^{-1/2}\right\| \right\} \geq \theta. \tag{38}$$

where $\delta W^N = \delta X^N \oplus \delta Z^N$. From this it evidently follows that $\|W^{-1/2}\delta W^N W^{-1/2}\| \geq \theta$.

The optimality condition for the least squares problem (17) is that

$$W\nabla\phi(W)W - \delta W^N \in W\left(\mathcal{L}^\perp \oplus \mathcal{L}\right)W.$$

From this one can see that

$$\begin{aligned}\bar{X} = (1/\rho)\left(Z^{-1}\delta Z^N Z^{-1} + Z^{-1}\right) &\in \mathcal{L} + D,\\ \bar{Z} = (1/\rho)\left(X^{-1}\delta X^N X^{-1} + X^{-1}\right) &\in \mathcal{L}^\perp + C,\end{aligned}$$

where $\rho = (n + \nu\sqrt{n})/\langle X, Z\rangle$. Assume that (38) is false, then both $Z^{-1/2}\delta Z^N Z^{-1/2} + I \in$ **Int** $\mathcal{P}$ and $X^{-1/2}\delta X^N X^{-1/2} + I \in$ **Int** $\mathcal{P}$, because $\theta < 1$. Therefore $\bar{W} = \bar{X} \oplus \bar{Z} \in$ **Int** $\mathcal{P}$ and both $\bar{X}$ and $\bar{Z}$ are strictly feasible. Moreover

$$\rho Z^{1/2}\bar{X}Z^{1/2} - I = Z^{-1/2}\delta Z^N Z^{-1/2} \quad \rho X^{1/2}\bar{Z}X^{1/2} - I = X^{-1/2}\delta X^N X^{-1/2},$$

which, from Lemma 2, implies

$$\rho\langle X, \bar{Z}\rangle + F(X) + F(\bar{Z}) \leq n + n\log\rho - \theta - \log(1-\theta) \tag{39}$$
$$\rho\langle \bar{X}, Z\rangle + F(\bar{X}) + F(Z) \leq n + n\log\rho - \theta - \log(1-\theta) \tag{40}$$

On the other hand, we also have (from the first part of Lemma 2)

$$\rho\langle \bar{X}, \bar{Z}\rangle + F(\bar{X}) + F(\bar{Z}) \geq n + n\log\rho.$$

Subtracting this from the sum of (39) and (40) we get

$$\rho\langle X, Z\rangle + F(X) + F(Z) \leq n + n\log\rho - 2\theta - 2\log(1-\theta)$$
$$\nu\sqrt{n} - \log\det XZ \leq n\log(n + \nu\sqrt{n}) - n\log\langle X, Z\rangle - 2\theta - 2\log(1-\theta). \tag{41}$$

Now $-2\theta - 2\log(1-\theta) < 1/6$ if $\theta \leq 0.35$ and therefore (41) would imply

$$\begin{aligned}
\psi(X,Z) &= n\log\langle X,Z\rangle - \log\det XZ - n\log n \\
&< n\log(1+\nu/\sqrt{n}) - \nu\sqrt{n} + 1/6 \\
&\leq n(\log(1+1/\sqrt{n}) - 1/\sqrt{n}) + 1/6 \\
&\leq -\frac{1}{2} + \frac{1}{3\sqrt{n}} + 1/6 \leq 0
\end{aligned}$$

where we made use of the fact that $\log(1+x) \leq x - \frac{1}{2}x^2 + \frac{1}{3}x^3$ for $x > -1$. The assumption that (38) is false therefore leads to a contradiction because $\psi(X,Z) \geq 0$ for all strictly feasible X, Z. □

A.3 Proof of Theorem 4

We will show that

$$\varphi(X - p\delta X, Z - q\delta Z) \leq \varphi(X,Z) - \delta.$$

This will prove the theorem because, as we have seen, for all feasible $\tilde{W}$, $\phi(\tilde{W}) \leq \varphi(\tilde{W})$, with equality for $\tilde{W} = W$. This follows from the fact that φ is obtained from ϕ by linearizing the concave term around W. For $s \in \mathbf{R}$ with $W - s\delta W > 0$, we have

$$\begin{aligned}
&\varphi(W - s\delta W) - \varphi(W) \\
&\quad= -\rho\langle Z \oplus X, s\delta W\rangle - \log\det(W - s\delta W) + \log\det W \\
&\quad= -s\langle \nabla\phi(W), \delta W\rangle + s\langle \nabla F(W), \delta W\rangle - \log\det(W - s\delta W) + \log\det W \\
&\quad= -s\langle \nabla\phi(W), \delta W\rangle - s\langle W^{-1}, \delta W\rangle - \log\det(W - s\delta W) + \log\det W \\
&\quad= -s\langle \nabla\phi(W), \delta W\rangle - s\mathbf{Tr}\,(W^{-1/2}\,\delta W W^{-1/2}) - \log\det(I - s\;W^{-1/2}\,\delta W\;W^{-1/2}).
\end{aligned}$$

Let $\mu_1, \ldots, \mu_{2n}$ be the eigenvalues of $W^{-1/2}\,\delta W\;W^{-1/2}$, and define μ as

$$\mu = (\sum_{i=1}^{2n} \mu_i^2)^{1/2} = \left\| W^{-1/2}\delta W W^{-1/2} \right\|.$$

Then from (15), we know that $\langle \nabla\phi(W), \delta W\rangle \geq \theta\mu$, and hence

$$\begin{aligned}
\varphi(W - s\delta W) - \varphi(W) &\leq -s\theta\mu - s\sum_{i=1}^{2n}\mu_i - \log\prod_{i=1}^{2n}(1 - s\mu_i) \\
&\leq -s\,(\theta+1)\mu - \log(1 - s\mu)
\end{aligned}$$

where we have used Lemma 1 (hence the inequality is only valid for for $-1/\mu < s < 1/\mu$). The upper bound is minimized for $s = \dfrac{\theta}{\mu(1+\theta)}$, which is precisely the value of p and q used in the theorem. We therefore find with this value of s,

$$\varphi(W - s\delta W) - \varphi(W) \leq -\theta + \log(1+\theta) = -\delta.$$

□