

# Bundle Method Sketching for Low Rank Semidefinite Programming

Lijun Ding

Benjamin Grimmer

Cornell University Ithaca, NY 14850, USA

LD446@CORNELL.EDU

BDG79@CORNELL.EDU

## Abstract

In this paper, we show that the bundle method can be applied to solve semidefinite programming problems with a low rank solution without ever constructing a full matrix. To accomplish this, we use recent results from randomly sketching matrix optimization problems and from the analysis of bundle methods. Under strong duality and strict complementarity of SDP, we achieve  $\tilde{O}(\frac{1}{\epsilon})$  convergence rates for both the primal and the dual sequences, and the algorithm proposed outputs a  $O(\sqrt{\epsilon})$  approximate solution  $\hat{X}$  (measured by distances) with a low rank representation with at most  $\tilde{O}(\frac{1}{\epsilon})$  many iterations.

## 1. Introduction

We consider solving semidefinite problems of the following form:

$$\begin{aligned} & \underset{X \in \mathbb{S}^n \subset \mathbb{R}^{n \times n}}{\text{maximize}} && \langle -C, X \rangle \\ & \text{subject to} && \mathcal{A}X = b, \quad X \succeq 0, \end{aligned} \tag{P}$$

where  $C \in \mathbb{S}^n \subset \mathbb{R}^{n \times n}$ ,  $\mathcal{A} : \mathbb{S}^n \rightarrow \mathbb{R}^m$  being linear, and  $b \in \mathbb{R}^m$ . Denote the solution set as  $\mathcal{X}_*$ . To accomplish the task of solving (P), we consider the dual problem:

$$\begin{aligned} & \underset{y \in \mathbb{R}^d}{\text{minimize}} && \langle -b, y \rangle \\ & \text{subject to} && \mathcal{A}^*y \preceq C, \end{aligned} \tag{D}$$

whose solution set is denoted as  $\mathcal{Y}_*$ . Then for all sufficiently large  $\alpha$ , e.g., larger than the trace of any solution  $X_* \in \mathcal{X}_*$  [5, Lemma 6.1], we can reformulate this as

$$\underset{y \in \mathbb{R}^d}{\text{minimize}} \quad F(y) = \langle -b, y \rangle - \alpha \min\{\lambda_{\min}(C - \mathcal{A}^*y), 0\}. \tag{1}$$

We propose applying the bundle method to solve this problem, which generates a sequence of dual solutions  $y_t$ . While the bundle method runs on the dual problem, a primal solution  $X_t$  can be constructed through a series of rank one updates corresponding to subgradients of  $F$ . However maintaining such a primal solution greatly increases memory costs. Fortunately, the primal problem enjoys having a low rank solution in many applications, e.g., matrix completion [11] and phase retrieval [4]. Also, without specifying the detailed structure the problem, there always exists a solution  $X_*$  to (P) with rank  $r_*$  satisfying  $\frac{r_*(r_*+1)}{2} \leq m$  [9]

To utilize the existence of such a low rank solution, we employ the matrix sketching methods introduced in [13]. The main idea is the following: the sketching method forms a linear sketch of the column and row spaces of the primal decision variable  $X$ , and then uses

the sketched column and row spaces to recover the primal decision variable. The recovered decision variable approximates the original  $X$  well if  $X$  is (approximately) low rank. Notably, we need not store the entire decision variable  $X$  at each iteration, but only the sketch. Hence the memory requirements of the algorithm are substantially reduced.

**Our Contributions.** Our proposed sketching bundle method produces a sequence of dual solutions  $y_k$  and a sequence of primal solutions  $X_t$  (which are sketched by low rank matrices  $\hat{X}_t$ ). This is done without ever needing to write down the solutions  $X_t$ , which can be a substantial boon to computational efficiency.

In particular, we consider problems satisfying the following pair of standard assumptions: (i) **Strong duality** holds, meaning that there is a solution pair  $(X_\star, y_\star) \in \mathcal{X}_\star \times \mathcal{Y}_\star$  satisfying

$$p_\star := \langle -C, X \rangle = \langle -b, y \rangle =: d_\star$$

and (ii) **Strict Complementarity** holds, meaning there is a solution pair  $(X_\star, y_\star)$  satisfying

$$\mathbf{Rank}(X_\star) + \mathbf{Rank}(C - \mathcal{A}^\star(y_\star)) = n.$$

Under these conditions, we show  $X_t$  and  $y_t$  converge in terms of primal and dual feasibility and optimality. In particular, all three of these quantities converge at a rate of  $O(1/t)$ .

**Theorem 1 (Primal-Dual Convergence)** *Suppose the sets  $\mathcal{X}_\star, \mathcal{Y}_\star$  are both compact, and that strong duality and a strict complementary condition holds. For any  $\epsilon > 0$ , Algorithm 1 with properly chosen parameters produces a solution pair  $X_t$  and  $y_t$  with*

$$t \leq \tilde{O}(1/\epsilon)$$

that satisfies

$$\begin{aligned} \text{approximate primal feasibility:} & \quad \|b - \mathcal{A}X_t\| \leq \epsilon, \quad X_t \succeq 0, \\ \text{approximate dual feasibility:} & \quad \lambda_{\min}(C - \mathcal{A}^\star y_t) \geq -\epsilon, \\ \text{approximate primal-dual optimality:} & \quad |\langle b, y_t \rangle - \langle C, X_t \rangle| \leq \epsilon \end{aligned}$$

Moreover, we show that assuming all of the minimizers in  $\mathcal{X}_\star$  are low rank, the sketched primal solutions  $\hat{X}_t$  converge to the set of minimizers at the following rate.

**Theorem 2 (Sketched Solution Convergence)** *Suppose the sets  $\mathcal{X}_\star, \mathcal{Y}_\star$  are both compact, strong duality and a strict complementary condition holds, and all solutions  $X_\star \in \mathcal{X}_\star$  have rank at most  $r$ . For any  $\epsilon > 0$ , Algorithm 1 with properly chosen parameters produces a sketched primal solution  $\hat{X}_t$  with*

$$t \leq \tilde{O}(1/\epsilon^2)$$

that satisfies  $\mathbb{E} \text{dist}_F(\hat{X}_t, \mathcal{X}_\star) \leq \epsilon$ .

## 2. Defining The Sketching Bundle Method

Our proposed proximal bundle method relies on an approximation of  $F$  given by  $\tilde{F}^t(y) = \max\{F(z_t) + \langle g_t, y - z_t \rangle, \bar{F}^t(y)\}$  where  $\bar{F}^t(y)$  is a convex combination of the lower bounds

from previous subgradients. Notice that a subgradient  $g_k$  of  $F$  at  $z_t$  can be computed as

$$v_t = \begin{cases} \text{MinEigenvector}(C - \mathcal{A}^*z_t) & \text{if } \lambda_{\min}(C - \mathcal{A}^*z_t) \leq 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

$$g_k = -b + \alpha \mathcal{A}v_t v_t^* \in \partial F(z_t). \quad (3)$$

Each iteration then computes a proximal step on this piecewise linear function given by

$$z_{t+1} \in \operatorname{argmin}_{y \in \mathbb{R}^d} \tilde{F}^t(y) + \frac{\rho}{2} \|y - y_t\|^2 \quad (4)$$

for some  $\rho > 0$ . The optimality condition of this subproblem ensures that for some  $\theta_t \in [0, 1]$   $0 = \theta_t g_t + (1 - \theta_t) (\nabla \bar{F}^t(z_t) + \rho(z_{t+1} - y_t))$ . Our aggregate lower bound  $\bar{F}^t(y)$  is updated to match this certifying subgradient  $\bar{F}^{t+1}(y) = \bar{F}^t(z_{t+1}) + \langle \theta_t g_t + (1 - \theta_t) \nabla \bar{F}^t(z_t), y - z_t \rangle$ . Then the following is an exact solution for the subproblem (4):

$$\theta_t = \min \left\{ 1, \frac{\rho(F(z_t) - \bar{F}^t(z_t))}{\|g_t - \nabla \bar{F}^t(z_t)\|^2} \right\} \quad (5)$$

$$z_{t+1} = y_t - \frac{1}{\rho} \left( \theta_t g_t + (1 - \theta_t) \nabla \bar{F}^t(z_t) \right). \quad (6)$$

If the decrease in value of  $F$  from  $y_t$  to  $z_{t+1}$  is at least  $\beta$  fraction of the decrease in value of  $\tilde{F}^t$  from  $y_t$  to  $z_{t+1}$ , then the bundle method sets  $y_{t+1} = z_{t+1}$  (called a descent step). Otherwise the method sets  $y_{t+1} = y_t$  (called a null step).

---

**Algorithm 1** Bundle Method with Cut Aggregation
 

---

**Input:**  $A: \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^d$ ,  $b \in \mathbb{R}^d$ ,  $C \in \mathbb{R}^{n \times n}$ ,  $y_1 \in \mathbb{R}^d$ ,  $\rho > 0$ ,  $\beta \in (0, 1)$ ,  $T \geq 0$ ,  $\alpha > 0$   
 Define  $\bar{F}^1(\cdot) = -\infty$ ,  $z_1 = y_1$ ,  $X_1 = 0$   
**for**  $t = 1, 2, \dots, T$  **do**  
     Compute  $v_t$  and  $g_t$  as in (2) and (3) *Compute One Subgradient*  
     Compute  $\theta_t$  and  $z_{t+1}$  as in (5) and (6) *Compute One Step*  
     **if**  $F(z_{t+1}) \leq F(y_t) - \beta(F(y_t) - \tilde{F}^k(z_{t+1}))$  **then**  $y_{t+1} = z_{t+1}$  *Take Descent Step*  
     **else**  $y_{t+1} = y_t$  **end if** *Take Null Step*  
     Set  $\bar{F}^{t+1}(y) = \bar{F}^t(z_{t+1}) + \langle \theta_t g_t + (1 - \theta_t) \nabla \bar{F}^t(z_t), y - z_t \rangle$  *Update Model of F*  
**end for**

---

**Extracting Primal Solutions Directly.** A solution to the primal problem (P) can be extracted from the sequence of subgradients  $g_k$  (as these describe the dual of the dual problem). Set our initial primal solution to be  $X_0 = 0$ . When each iteration updates the model  $\bar{F}^t$ , we could update our primal variable similarly:

$$X_{t+1} = \theta_t \alpha v_t v_t^* + (1 - \theta_t) X_t. \quad (7)$$

As stated in Theorem 1, these primal solutions  $X_t$  converge to optimality and feasibility at a rate of  $O(1/t)$ . Alas this approach requires us to compute and store the full  $n \times n$  matrix  $X_t$  at each iteration. Assuming every solution is low rank, an ideal method would only require  $O(nr)$  memory. The following section shows matrix sketching can accomplish this.

**Extracting Primal Solutions Via Matrix Sketching.** Here we describe how the matrix sketching method of [13] can be used to store an approximation of our primal solution  $X_t \in \mathbb{S}_+^n$ . First, we draw two matrices with independent standard normal entries

$$\begin{aligned}\Psi &\in \mathbb{R}^{n \times k} \quad \text{with } k = 2r + 1; \\ \Phi &\in \mathbb{R}^{l \times n} \quad \text{with } l = 4r + 3;\end{aligned}$$

Here  $r$  is chosen by the user. It either represents the estimate of the true rank of the primal solution or the user’s computational budget in dealing with larges matrices.

We use  $Y_t^C$  and  $Y_t^R$  to capture the column space and the row space of  $X_t$ :

$$Y_t^C = X_t \Psi \in \mathbb{R}^{n \times k}, \quad Y_t^R = \Phi X_t \in \mathbb{R}^{l \times n}. \quad (8)$$

Hence we initially have  $Y_0^C = 0$  and  $Y_0^R = 0$ . Notice Algorithm 1 does not observe the matrix  $X_t$  directly. Rather, it observes a stream of rank one updates

$$X_{t+1} = \theta_t \alpha v_t v_t^* + (1 - \theta_t) X_t.$$

In this setting,  $Y_{t+1}^C$  and  $Y_{t+1}^R$  can be directly computed as

$$Y_{t+1}^C = \theta_t \alpha v_t (v_t^* \Psi) + (1 - \theta_t) Y_t^C \in \mathbb{R}^{n \times k}, \quad (9)$$

$$Y_{t+1}^R = \theta_t \alpha (\Psi v_t) v_t^* + (1 - \theta_t) Y_t^R \in \mathbb{R}^{l \times n}. \quad (10)$$

This observation allows us to form the sketch  $Y_t^C$  and  $Y_t^R$  from the stream of updates.

We then reconstruct  $X_t$  and get the reconstructed matrix  $\hat{X}_t$  by

$$Y_t^C = Q_t R_t, \quad B_t = (\Phi Q_t)^\dagger Y_t^R, \quad \hat{X}_t = Q_t [B_t]_r, \quad (11)$$

where  $Q_t R_t$  is the  $QR$  factorization of  $Y_t^C$  and  $[\cdot]_r$  returns the best rank  $r$  approximation in Frobenius norm. Specifically, the best rank  $r$  approximation of a matrix  $Z$  is  $U \Sigma V^*$ , where  $U$  and  $V$  are right and left singular vectors corresponding to the  $r$  largest singular values of  $Z$  and  $\Sigma$  is a diagonal matrix with  $r$  largest singular values of  $Z$ . In actual implementation, we may only produce the factors  $(QU, \Sigma, V)$  defining  $\hat{X}_T$  in the end instead reconstructing  $\hat{X}_t$  in every iteration.

Hence the expensive primal update (7) can be replaced by much more efficient operations (9) and (10). Then a low rank approximation of  $X_t$  can be computed by (11).

We remark that the reconstructed matrix  $\hat{X}_t$  is not necessarily positive definite. However, this suffices for the purpose of finding a matrices close to  $X_t$ . More sophisticated procedure is available for producing a positive semidefinite approximation of  $X_t$  [13, Section 7.3].

### 3. Numerical Experiments

In this section, we demonstrate Algorithm 1 equipped with the sketching procedure does solve problem instances in four important problem classes: (1) Generalized eigenvalue [3, Section 5.1], (2)  $\mathbb{Z}_2$  synchronization [2], (3) Max-Cut [7], and (4) matrix completion [11]. For all experiments, we set  $\beta = \frac{1}{4}$ ,  $\rho = 1$ , and  $\alpha = 2 \langle I, X_\star \rangle$  where  $X_\star$  is computed via MOSEK [8] We present the results of convergences in Figure 1 for the following problem instances (results for other problems instances are similar for each problem class): Let  $\mathcal{D}$  be the distribution of symmetric matrices in  $\mathbb{S}^{800}$  with upper triangular part (including the diagonal) being independent standard Gaussians.

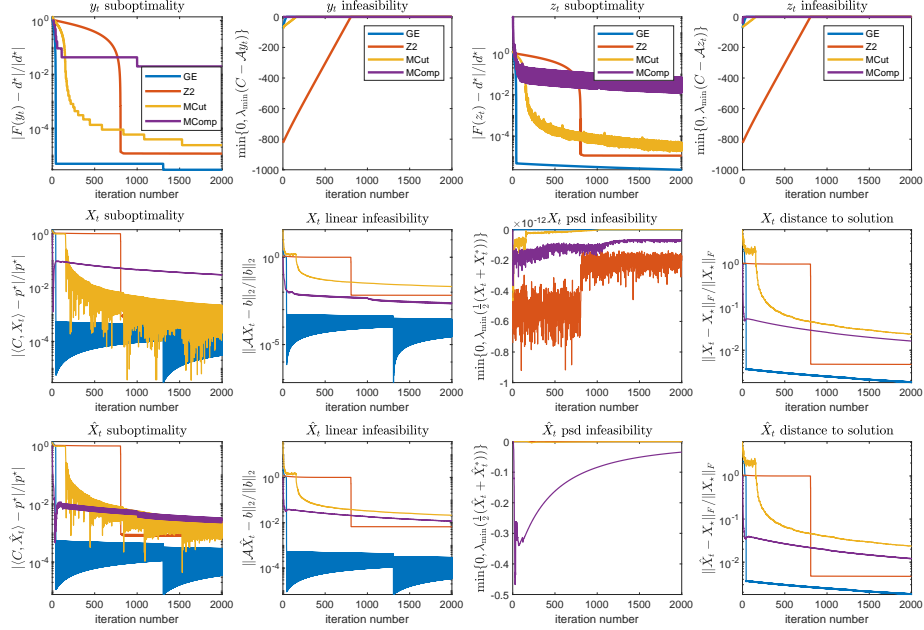


Figure 1: The upper, middle, and lower plots give the different convergence measures (as indicated on the title) for the dual iterates  $z_t$  and  $y_t$ , the primal iterates  $X_k$ , and the sketching iterate  $\hat{X}_t$  respectively. The color corresponds to different problem instances as indicated in the legend of the upper plots. Due to numerical error in updating  $X_t$  and non-asymmetry of  $\hat{X}_t$ , we compute the minimum eigenvalue of their symmetrized version.  $\lambda_{\min}(\frac{1}{2}(X_t + X_t^T))$  appears to be negative but very small due to numerical error.

1. Generalized eigenvalue (**GE**):  $C \sim D$ ,  $\mathcal{A}(X) = \langle B, X \rangle$  for any  $X$ , where  $B = I + \frac{1}{n}WW^*$  and  $W \sim D$  and is independent of  $C$ , and  $b = 1 \in \mathbb{R}$ .
2.  $\mathbb{Z}_2$  synchronization (**Z2**):  $C = J + \frac{1}{2\sqrt{5}}W$  where  $J \in \mathbb{S}^{800}$  is the all one matrix and  $W \sim D$ ,  $\mathcal{A} = \mathbf{diag}$ , and  $b \in \mathbb{R}^{800}$  is the all one vector.
3. Max-Cut(**MCut**):  $C = -L$  where  $L$  is the Laplacian matrix of the G1 graph [1] with 800 vertices,  $\mathcal{A}$  and  $b$  are the same as  $\mathbb{Z}_2$  synchronization
4. Matrix Completion(**MComp**): A random rank 3 matrix  $X^{\natural} \in \mathbb{S}^{400}$  is generated. The index set  $\Omega \subset \{1, \dots, 400\} \times \{1, \dots, 400\} =: [400]^2$  is generated in a way that each  $(i, j) \in [400]^2$  is in  $\Omega$  with probability 0.2 independently from anything else. Set  $C = I \in \mathbb{S}^{800}$ . The linear constraint is  $X_{n+i, n+j} = X_{i, j}^{\natural}$  for each  $(i, j) \in \Omega$ . So  $[\mathcal{A}(X)]_{i, j} = X_{i+n, j+n}$  and  $b_{i, j} = X_{i, j}^{\natural}$  for each  $(i, j) \in \Omega$ .

In general, Figure 1 shows the method is effective. We note the convergence is quicker when the problems have rank 1 optimal solutions (**GE** and **Z2**) comparing to problems with higher rank optimal solutions (**MCut** and **MComp**).

## ACKNOWLEDGMENTS

L. Ding was partially supported by the National Science Foundation CRII award 1657420 and grant 170482. B. Grimmer was supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1650441.

## References

- [1] The university of florida sparse matrix collection: Gset group. URL <https://www.cise.ufl.edu/research/sparse/matrices/Gset/index.html>.
- [2] Afonso S Bandeira. Random laplacian matrices and convex relaxations. *Foundations of Computational Mathematics*, 18(2):345–379, 2018.
- [3] Nicolas Boumal, Vladislav Voroninski, and Afonso S Bandeira. Deterministic guarantees for burer-monteiro factorizations of smooth semidefinite programs. *Communications on Pure and Applied Mathematics*, 2018.
- [4] Emmanuel J Candes, Thomas Strohmer, and Vladislav Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- [5] Lijun Ding, Alp Yurtsever, Volkan Cevher, Joel A Tropp, and Madeleine Udell. An optimal-storage approach to semidefinite programming using approximate complementarity. *arXiv preprint arXiv:1902.03373*, 2019.
- [6] Yu Du and Andrzej Ruszczyński. Rate of convergence of the bundle method. *J. Optim. Theory Appl.*, 173(3):908–922, June 2017. ISSN 0022-3239. doi: 10.1007/s10957-017-1108-1. URL <https://doi.org/10.1007/s10957-017-1108-1>.
- [7] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- [8] APS Mosek. The mosek optimization software. *Online at http://www.mosek.com*, 54(2-1):5, 2010.
- [9] Gábor Pataki. On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Mathematics of operations research*, 23(2):339–358, 1998.
- [10] Andrzej P Ruszczyński and Andrzej Ruszczyński. *Nonlinear optimization*, volume 13. Princeton university press, 2006.
- [11] Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In *International Conference on Computational Learning Theory*, pages 545–560. Springer, 2005.
- [12] Jos F Sturm. Error bounds for linear matrix inequalities. *SIAM Journal on Optimization*, 10(4):1228–1248, 2000.
- [13] Joel A Tropp, Alp Yurtsever, Madeleine Udell, and Volkan Cevher. Practical sketching algorithms for low-rank matrix approximation. *SIAM Journal on Matrix Analysis and Applications*, 38(4): 1454–1485, 2017.
- [14] Joel A Tropp, Alp Yurtsever, Madeleine Udell, and Volkan Cevher. Randomized single-view algorithms for low-rank matrix approximation. 2017.

## Appendix A. Proofs of Convergence Guarantees

### A.1. Auxiliary lemmas

**Lemma 3** (*Sketching Guarantee*) [14, Theorem 5.1] Fix a target rank  $r$ . Let  $X$  be a matrix, and let  $(Y, W)$  be a sketch of  $X$  of the form (8). The procedure (11) yields a rank- $r$  matrix  $\hat{X}$  with

$$\mathbb{E} \left\| X - \hat{X} \right\|_F \leq 3\sqrt{2} \|X - [X]_r\|_F.$$

Here  $\|\cdot\|_F$  is the Frobenius norm. Similar operator 2-norm bounds hold with high probability.

**Lemma 4 (Compact sublevel set)** If a convex lower semicontinuous function  $f(x) : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$  has a compact nonempty solution set, then all of its sublevel set is compact.

**Proof** Suppose for some  $L \in \mathbb{R}$ , then the closed sublevel set  $S_L = \{x \mid f(x) \leq L\}$  is unbounded. Then there is a unit direction vector  $\gamma \in \mathbb{R}$  such that for all  $x \in S_L, \alpha \geq 0$ ,  $x + \alpha\gamma \in S_L$ . This in particular violates the fact the solution set is bounded and the proof is completed.  $\blacksquare$

**Lemma 5 (Quadratic Growth)** [12, Section 4] If the solution sets  $\mathcal{X}_\star$  and  $\mathcal{Y}_\star$  are compact and strict complementarity holds, then for any fixed  $\epsilon > 0$ , there are some  $\gamma_1, \gamma_2$  such that for all  $y$  with  $F(y) \leq \langle -b, y_\star \rangle + \epsilon$  with  $\alpha = 2 \sup_{X_\star \in \mathcal{X}_\star} \text{trace}(X_\star)$ , and all  $X \succeq 0$  with  $|\langle C, X \rangle - \langle C, X_\star \rangle| \leq \epsilon$  and  $\|\mathcal{A}X - b\|_2 \leq \epsilon$ :

$$\text{dist}_F^2(y, \mathcal{Y}_\star) \leq \gamma_1 (F(y) - F(y_\star)), \quad \text{dist}_F^2(X, \mathcal{X}_\star) \leq \gamma_2 (|\langle C, X \rangle - \langle C, X_\star \rangle| + \|\mathcal{A}X - b\|_2).$$

**Proof** Using the proof of Lemma 7, we find that  $\lambda_{\min}(C - \mathcal{A}^*(y)) \geq -\frac{\epsilon}{\sup_{X_\star \in \mathcal{X}_\star} \text{trace } X_\star}$ . Thus  $\langle -b, y \rangle \leq \langle -b, y_\star \rangle + 3\epsilon$ . The result in [12, Section 4] requires the set  $S_1 = \{y \mid F(y) \leq \langle -b, y_\star \rangle + \epsilon\}$ , and the set  $S_2 = \{X \mid X \succeq 0, |\langle C, X \rangle - \langle C, X_\star \rangle| \leq \epsilon, \text{ and } \|\mathcal{A}X - b\|_2 \leq \epsilon\}$  being compact. Using [10, Theorem 7.21], the optimization problem  $\min_{X \succeq 0} g(X) := \langle C, X \rangle + \gamma \|\mathcal{A}X - b\|_2$  has the same solution set as the primal SDP (P) for some large  $\gamma > 0$ . Thus the compactness of the set  $S_1$  and  $S_2$  is ensured by Lemma 4, and the proof is completed.  $\blacksquare$

### A.2. Proof of Theorem 1

Let  $D_{\mathcal{X}_\star} \geq \sup_{X_\star \in \mathcal{X}_\star} \text{trace}(X_\star)$  and  $D_y \geq \sup_{F(y) \leq F(y_0)} \|y\|_2$ . Then we set the bundle method's parameters to be  $\alpha = 2D_{\mathcal{X}_\star}$ ,  $\beta = 1/2$ , and  $\rho = \epsilon/D_y^2$ . We recall the inner product for matrices is the trace inner product and is the dot product for the vectors.

In the following three lemmas, we prove bounds on primal feasibility, dual feasibility, and optimality in terms of  $F(y_t) - F(y^*)$ . From this, we can conclude these quantities converge at the claimed rate since the Du and Ruzcynskii [6] recently showed the bundle method has  $F(y_t) - F(y^*)$  converge at a  $\tilde{O}(1/\epsilon)$  rate.

**Lemma 6 (Primal Feasibility)** *At every descent step  $t$ , we have approximate primal feasibility*

$$\begin{aligned} X_{t+1} &\succeq 0, \\ \|b - \mathcal{A}X_{t+1}\| &\leq \sqrt{\frac{2\epsilon(F(y_t) - F(y^*))}{\beta D_y^2}}. \end{aligned}$$

**Proof** Noting that  $X_{t+1}$  is built out of a convex combination of the rank matrices  $v_t v_t^T$ , it is immediate that it is always a positive semidefinite matrix.

The definition of  $X_t$  immediately gives the following alternative characterization of  $\bar{F}^{t+1}$ ,

$$\bar{F}^{t+1}(y) = -\langle C, X_{t+1} \rangle - \langle b - \mathcal{A}X_{t+1}, y \rangle.$$

Since we constructed  $\bar{F}^{t+1}$  to correspond to the first-order optimality condition of the subproblem (4), we have

$$0 = \nabla \bar{F}^{t+1}(y_{t+1}) + \rho(z_{t+1} - y_t).$$

Hence  $\|b - \mathcal{A}X_{t+1}\|^2 = \rho^2 \|y_{t+1} - y_t\|^2$ . The distance traveled during any descent step can be bounded by the objective value gap as

$$\frac{\rho}{2} \|y_{t+1} - y_t\|^2 \leq F(y_t) - \tilde{F}^t(y_{t+1}) \leq \frac{F(y_t) - F(y_{t+1})}{\beta} \leq \frac{F(y_t) - F(y_*)}{\beta}$$

where the first inequality uses the fact that  $z_{t+1}$  minimizes  $\tilde{F}^t(\cdot) + \frac{\rho}{2} \|\cdot - y_t\|^2$  and the second inequality uses the definition of a descent step. Combining this with our feasibility bound shows

$$\|b - \mathcal{A}X_{t+1}\|^2 \leq \frac{2\rho(F(y_t) - F(y_*))}{\beta}.$$

Then our choice of  $\rho$  completes the proof. ■

**Lemma 7 (Dual Feasibility)** *At every descent step  $t$ , we have approximate dual feasibility*

$$\lambda_{\min}(C - \mathcal{A}^* y_{t+1}) \geq \frac{-(F(y_t) - F(y_*))}{D_{X_*}}.$$

**Proof** Standard strong duality and exact penalization arguments show for any  $X_* \in \mathcal{X}_*$ ,

$$\begin{aligned} \langle b, y_{t+1} - y_* \rangle &= \langle \mathcal{A}X_*, y_{t+1} \rangle - \langle C, X_* \rangle \\ &\leq \langle X_*, \mathcal{A}^*(y_{t+1} - C) \rangle \\ &\leq -\text{trace}(X_*) \min\{\lambda_{\min}(C - \mathcal{A}^* y_{t+1}), 0\}. \end{aligned}$$

Recalling our assumption that  $\alpha \geq 2D \geq 2\text{trace}(X_*)$  yields the claimed feasibility bound

$$\begin{aligned} F(y_t) - F(y_*) &\geq F(y_{t+1}) - F(y_*) = \langle -b, y_{t+1} - y_* \rangle - \alpha \min\{\lambda_{\min}(C - \mathcal{A}^* y_{t+1}), 0\} \\ &\geq -\text{trace}(X_*) \min\{\lambda_{\min}(C - \mathcal{A}^* y_{t+1}), 0\}. \end{aligned}$$
■



**Lemma 8 (Primal-Dual Optimality)** *At every descent step  $t$ , we have approximate primal-dual optimality bounded above by*

$$\langle b, y_{t+1} \rangle - \langle C, X_{t+1} \rangle \leq \frac{\alpha}{D_{\mathcal{X}_*}} (F(y_t) - F(y_*)) + \sqrt{\frac{2\epsilon(F(y_t) - F(y_*))}{\beta}}.$$

and below by

$$\langle b, y_{t+1} \rangle - \langle C, X_{t+1} \rangle \geq -\frac{1-\beta}{\beta} (F(y_t) - F(y_*)) - \sqrt{\frac{2\epsilon(F(y_t) - F(y_*))}{\beta}}.$$

**Proof** The standard duality analysis shows the primal-dual objective gap equals

$$\begin{aligned} \langle b, y_{t+1} \rangle - \langle C, X_{t+1} \rangle &= \langle \mathcal{A}X_{t+1}, y_{t+1} \rangle - \langle C, X_{t+1} \rangle + \langle b - \mathcal{A}X_{t+1}, y_{t+1} \rangle \\ &= \langle X_{t+1}, \mathcal{A}^*y_{t+1} - C \rangle + \langle b - \mathcal{A}X_{t+1}, y_{t+1} \rangle. \end{aligned}$$

Notice that the second term here is bounded above and below as

$$|\langle b - \mathcal{A}X_{t+1}, y_{t+1} \rangle| \leq \sqrt{\frac{2\epsilon(F(y_t) - F(y_*))}{\beta D_y^2}} \|y_{t+1}\|_2 \leq \sqrt{\frac{2\epsilon(F(y_t) - F(y_*))}{\beta}}$$

by Lemma 6. Hence we only need to show that the first term also approaches zero (that is, we approach holding complementary slackness).

An upper bound on this inner product follows from Lemma 7 as

$$\langle X_{t+1}, \mathcal{A}^*y_{t+1} - C \rangle \leq -\text{trace}(X_{t+1}) \min\{\lambda_{\min}(C - \mathcal{A}^*y_{t+1}), 0\} \leq \frac{\text{trace}(X_{t+1})(F(y_t) - F(y_*))}{D_{\mathcal{X}_*}}.$$

Hence

$$\langle b, y_{t+1} \rangle - \langle C, X_{t+1} \rangle \leq \frac{\alpha}{D_{\mathcal{X}_*}} (F(y_t) - F(y_*)) + \sqrt{\frac{2\epsilon(F(y_t) - F(y_*))}{\beta}}.$$

A lower bound on this inner product follows as

$$\begin{aligned} \frac{1-\beta}{\beta} (F(y_t) - F(y_{t+1})) &\geq F(y_{t+1}) - \tilde{F}^t(y_{t+1}) \\ &= F(y_{t+1}) - \bar{F}^{t+1}(y_{t+1}) \\ &= -\alpha \min\{\lambda_{\min}(C - \mathcal{A}^*y_{t+1}), 0\} + \langle C, X_{t+1} \rangle - \langle \mathcal{A}X_{t+1}, y_{t+1} \rangle \\ &\geq \langle X_{t+1}, C - \mathcal{A}^*y_{t+1} \rangle, \end{aligned}$$

where the first inequality follows from the definition of a descent step. Hence

$$\langle b, y_{t+1} \rangle - \langle C, X_{t+1} \rangle \geq -\frac{1-\beta}{\beta} (F(y_t) - F(y_*)) - \sqrt{\frac{2\epsilon(F(y_t) - F(y_*))}{\beta}}.$$

■

### A.3. Proof of Theorem 2

Using triangle inequality, we see that

$$\mathbb{E} \operatorname{dist}_{\mathbb{F}}(\hat{X}_t, \mathcal{X}_\star) \leq \mathbb{E} \left\| \hat{X}_t - X_t \right\|_{\mathbb{F}} + \operatorname{dist}_{\mathbb{F}}(X_t, \mathcal{X}_\star). \quad (12)$$

The first term of (12) is bounded by

$$\mathbb{E} \left\| \hat{X}_t - X_t \right\|_{\mathbb{F}} \stackrel{(a)}{\leq} 3\sqrt{2} \|X_t - [X_t]_r\|_{\mathbb{F}} \stackrel{(b)}{\leq} 3\sqrt{2} \operatorname{dist}_{\mathbb{F}}(X_t, \mathcal{X}_\star),$$

where step (a) is due to Lemma 3, and step (b) is because  $X_\star \in \mathcal{X}_\star$  all has rank less than or equal to  $r$  and  $[X_t]_r$  is the best rank  $r$  approximation of  $X_t$  in terms of Frobenius norm.

Combining the above inequalities, we see that

$$\mathbb{E} \operatorname{dist}_{\mathbb{F}}(\hat{X}_t, \mathcal{X}_\star) \leq (1 + 3\sqrt{2}) \operatorname{dist}_{\mathbb{F}}(X_t, \mathcal{X}_\star). \quad (13)$$

Our task now is to have an estimate the rates convergence of  $\operatorname{dist}_{\mathbb{F}}(X_t, \mathcal{X}_\star)$ . Denote the shorthand  $g(X) = |\langle C, X \rangle - \langle C, X_\star \rangle| + \gamma \|AX - b\|_2$ . Using Lemma 5, there are some  $\gamma$  such that for all  $X$  with  $g(X) \leq g(X_\star) + 1$  satisfy

$$\operatorname{dist}_{\mathbb{F}}^2(X, \mathcal{X}_\star) \leq \gamma(g(X) - g(X_\star)). \quad (14)$$

Then the theorem follows from Theorem 1 as  $g(X_t) - g(X_\star)$  converges at a rate of  $\tilde{O}(1/\epsilon)$ .