

## MIT Open Access Articles

*Convergence rate of block-coordinate maximization  
Burer–Monteiro method for solving large SDPs*

The MIT Faculty has made this article openly available. *Please share*  
how this access benefits you. Your story matters.

**Citation:** Erdogdu, Murat A, Ozdaglar, Asuman, Parrilo, Pablo A and Vanli, Nuri Denizcan. 2021. "Convergence rate of block-coordinate maximization Burer–Monteiro method for solving large SDPs." Mathematical Programming.

**As Published:** 10.1007/S10107-021-01686-3

**Publisher:** Springer Science and Business Media LLC

**Persistent URL:** <https://hdl.handle.net/1721.1/143826>

**Version:** Original manuscript: author's manuscript prior to formal peer review

**Terms of use:** Creative Commons Attribution-Noncommercial-Share Alike



# Convergence Rate of Block-Coordinate Maximization Burer-Monteiro Method for Solving Large SDPs

Murat A. Erdogdu\*, Asuman Ozdaglar†, Pablo A. Parrilo†, N. Denizcan Vanli†

November 27, 2019

## Abstract

Semidefinite programming (SDP) with diagonal constraints arise in many optimization problems, such as Max-Cut, community detection and group synchronization. Although SDPs can be solved to arbitrary precision in polynomial time, generic convex solvers do not scale well with the dimension of the problem. In order to address this issue, Burer and Monteiro [BM03] proposed to reduce the dimension of the problem by appealing to a low-rank factorization and solve the subsequent non-convex problem instead. In this paper, we present coordinate ascent based methods to solve this non-convex problem with provable convergence guarantees. More specifically, we prove that the block-coordinate maximization algorithm applied to the non-convex Burer-Monteiro method globally converges to a first-order stationary point with a sublinear rate without any assumptions on the problem. We further show that this algorithm converges linearly around a local maximum provided that the objective function exhibits quadratic decay. We establish that this condition generically holds when the rank of the factorization is sufficiently large. Furthermore, incorporating Lanczos method to the block-coordinate maximization, we propose an algorithm that is guaranteed to return a solution that provides  $1 - \mathcal{O}(1/r)$  approximation to the original SDP without any assumptions, where  $r$  is the rank of the factorization. This approximation ratio is known to be optimal (up to constants) under the unique games conjecture, and we can explicitly quantify the number of iterations to obtain such a solution.

## 1 Introduction

A variety of problems in statistical estimation and machine learning require solving a combinatorial optimization problem, which are often intractable [VB96]. Semidefinite programs (SDP) are commonly used as convex relaxations for these problems, providing efficient algorithms with approximate optimality [Par03]. A generic SDP in this framework can be written as

$$\begin{aligned} & \text{maximize} && \langle \mathbf{A}, \mathbf{X} \rangle && \text{(CVX)} \\ & \text{subject to} && X_{ii} = 1, \text{ for } i \in [n], \\ & && \mathbf{X} \succeq 0, \end{aligned}$$

where  $\mathbf{A}, \mathbf{X} \in \text{Sym}_n$  (real symmetric matrices of size  $n \times n$ ) and  $[n] = \{1, 2, \dots, n\}$ . This problem appears as a convex relaxation to the celebrated Max-Cut problem [GW95], graphical model inference [EDM17], community detection problems [BBV16], and group synchronization [MMMO17].

---

\*University of Toronto, Department of Computer Science and Department of Statistical Sciences, erdogdu@cs.toronto.edu

†Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science {asuman,parrilo,denizcan}@mit.edu

Although SDPs serve as reliable relaxations to many combinatorial problems, the resulting convex problem is still computationally challenging. Interior point methods can solve SDPs to arbitrary accuracy in polynomial-time, but they do not scale well with the problem dimension  $n$ . A popular approach to remedy these limitations is to introduce a low-rank factorization  $\mathbf{X} = \boldsymbol{\sigma}\boldsymbol{\sigma}^\top$ , where  $\boldsymbol{\sigma} \in \mathbb{R}^{n \times r}$  with  $r$  denoting the rank. This reformulation removes the positive semidefinite cone constraint in (CVX) since  $\mathbf{X} = \boldsymbol{\sigma}\boldsymbol{\sigma}^\top$  is guaranteed to be a positive semidefinite matrix, and choosing  $r \ll n$  provides computational efficiency as well as storage benefits. This method is often referred to as Burer-Monteiro approach [BM03]. Denoting  $i$ -th row of  $\boldsymbol{\sigma}$  by  $\sigma_i$ , i.e.,  $\boldsymbol{\sigma} = [\sigma_1, \sigma_2, \dots, \sigma_n]^\top$ , the resulting non-convex problem can be written as follows

$$\begin{aligned} & \text{maximize} && \langle \mathbf{A}, \boldsymbol{\sigma}\boldsymbol{\sigma}^\top \rangle && \text{(Non-CVX)} \\ & \text{subject to} && \|\sigma_i\| = 1, \text{ for } i \in [n]. \end{aligned}$$

In the original Burer-Monteiro approach [BM03], the authors proposed to use an augmented Lagrangian method for a general form SDP. However, it has been recently observed that feasible methods (such as block-coordinate maximization [JMRT16, WCK17], Riemannian gradient [JMRT16, MMMO17] and Riemannian trust-region methods [ABG07, BVB16, JBAS10]) provide empirically faster rates since feasibility can be efficiently guaranteed via projection onto the Cartesian product of spheres. Despite overwhelming empirical evidence [JMRT16, MMMO17, WCK17], convergence properties of these feasible methods are not well-understood (except the Riemannian trust-region method, for which a sublinear convergence rate is shown in [BVB16] and a local superlinear convergence is shown in [ABG07] with no rate estimate). Among these methods, block-coordinate maximization and Riemannian gradient ascent are simpler to implement and have computational complexity of  $\mathcal{O}(nr)$  and  $\mathcal{O}(n^2r)$ , respectively, whereas Riemannian trust-region requires to solve the trust-region subproblem at each iteration, which is usually solved iteratively using the Lanczos method in a few iterations, whose per iteration requires  $\mathcal{O}(n^2r)$  arithmetic operations. Furthermore, block-coordinate maximization does not have any step size or tuning parameters, unlike Riemannian gradient ascent and Riemannian trust-region methods. Empirical studies further motivate the use of block-coordinate maximization by presenting superior performance compared to existing methods on large-scale problems, often with linear convergence [JMRT16]. In this paper, we provide the first local and global convergence rate guarantees for the block-coordinate maximization method (applied to Burer-Monteiro approach) in the literature, which are consistent with the empirical performance of the algorithm. Our contributions can be summarized as follows:

- We establish the global sublinear convergence of the block-coordinate maximization algorithm applied to (Non-CVX) without any assumptions on the cost matrix  $\mathbf{A}$ .
- We show that this algorithm enjoys a linear rate around a neighborhood of any local maximum when the objective function satisfies the quadratic decay assumption.
- We establish that the quadratic decay condition that leads to local linear convergence generically holds when the rank of the factorization satisfies  $r \geq \sqrt{2n}$ .
- Incorporating Lanczos methods into the block-coordinate maximization procedure, we propose an algorithm that returns an approximate second-order stationary point of (Non-CVX). By choosing the rank of the factorization sufficiently large and selecting the parameters of the algorithm according to the cost matrix  $\mathbf{A}$ , we show that the solution returned by this algorithm is not only an approximate local maximum to (Non-CVX), but also provides  $1 - \mathcal{O}(1/r)$  approximation to (CVX). We highlight that this approximation ratio is optimal under the unique games conjecture.

- We validate our theoretical results via numerical examples and compare the performance of the block-coordinate maximization algorithm with various manifold optimization methods to demonstrate its performance.

## 1.1 Related Work

There are numerous papers that analyze the landscape of the solution space of **(Non-CVX)**. In particular, it is known that **(CVX)** admits an optimal solution of rank  $r$  such that  $r(r+1)/2 \leq n$  [Bar95, Pat98]. Using this observation, it has been shown in [BM03, BM05, JBAS10] that when  $r \geq \sqrt{2n}$ , if  $\sigma$  is a rank deficient second-order stationary point of **(Non-CVX)**, then  $\sigma$  is a global maximum for **(Non-CVX)** and  $\mathbf{X} = \sigma\sigma^\top$  is a global maximum for **(CVX)**. The recent paper [BVB18] showed that when  $r \geq \sqrt{2n}$ , for almost all  $\mathbf{A}$ , every  $\sigma$  that is a first-order stationary point is rank deficient. For arbitrary rank  $r$ , it is shown that all local maxima are within a  $n\|\mathbf{A}\|_2/\sqrt{r}$  gap from the optimum of **(CVX)** [Mon16], and any  $\varepsilon$ -approximate concave point is within a  $\text{Rg}(\text{Non-CVX})/(r-1) + n\varepsilon/2$  gap from the optimum of **(CVX)** [MMMO17], where  $\text{Rg}(\text{Non-CVX})$  is the range of the problem **(Non-CVX)**, i.e., the difference between the maximum and the minimum values of the objective in **(Non-CVX)**.

Javanmard *et al.* [JMRT16] showed that when applied to solve **(Non-CVX)**, Riemannian gradient ascent and block-coordinate maximization methods provide excellent numerical results, yet no convergence guarantee is provided. Similar experimental results are also observed in [WCK17] for the block-coordinate maximization algorithm and in [MMMO17] for the Riemannian gradient ascent algorithm. Concurrent to this work, in [WCK17], the authors analyzed the convergence of the deterministic block-coordinate maximization algorithm. In particular, they showed that the deterministic block-coordinate maximization algorithm is asymptotically convergent (see [WCK17, Theorem 3.2]) and enjoys a local linear convergence with no explicit rate estimates (see [WCK17, Theorem 3.5]). They also proved that the deterministic block-coordinate maximization approach converges to a local maximum generically under random initialization using the center-stable manifold theorem similar to [LSJR16]. These results hold under the assumption that the iterates generated by the algorithm satisfy a certain condition that is seemingly impossible to verify without actually running the algorithm. To alleviate this issue, the authors suggested using a coordinate ascent method with a sufficiently small step size, for which the aforementioned convergence results hold without this precarious assumption. In [BAC16], the authors provided a global sublinear convergence rate for the Riemannian trust-region method for general non-convex problems and these results have been used in [BVB16, MMMO17] for the non-convex Burer-Monteiro approach. Augmented Lagrangian methods have been proposed to solve **(Non-CVX)** as well [BM03, BM05], however these methods do not benefit from separability of the manifold constraints, and hence are usually slower [BVB18].

There also exist methods that solve **(CVX)** by exploiting its special structure [AHK05, GH11, KL96, Ste10]. In particular, [KL96] reduces **(CVX)** to a sequence of approximate eigenpair computations that is efficiently solved using the power method. In [AHK05, Ste10], matrix multiplicative weights algorithm is used to approximately solve **(CVX)**, and these ideas are extended in [GH11] using sketching techniques [TYUC17]. However, these methods require constructing  $\mathbf{X} \in \text{Sym}_n$  explicitly, which is prohibitive when  $n$  goes beyond a few thousands, whereas the Burer-Monteiro approach we consider in this paper easily scales to very large instances as the low-rank factorization decreases the dimension of the problem from  $\mathcal{O}(n^2)$  to  $\mathcal{O}(nr)$  with  $r \ll n$ . For time complexity comparison between these methods that are based on Lagrangian relaxation and the Burer-Monteiro approach in this paper, we refer to Corollary 11.

## 1.2 Notations and Preliminaries

Throughout the paper, matrices are denoted with a boldface font, and all vectors are column vectors. The superscripts are used to denote iteration counters, i.e.,  $\boldsymbol{\sigma}^k$  denotes the value of  $\boldsymbol{\sigma}$  at iteration  $k$ . For a vector  $g$ ,  $\|g\|$  represents its Euclidean norm. For matrices  $\mathbf{A}, \mathbf{B}$ , we write  $\langle \mathbf{A}, \mathbf{B} \rangle = \text{trace}(\mathbf{A}\mathbf{B}^\top)$  for the inner product associated to the Frobenius norm  $\|\mathbf{A}\|_F = \sqrt{\langle \mathbf{A}, \mathbf{A} \rangle}$ .  $A_{ij}$  represents the entry at the  $i$ -th row and  $j$ -th column of  $\mathbf{A}$ ,  $A_i$  represents its  $i$ -th row as a column vector, and  $\|\mathbf{A}\|_1 = \max_{1 \leq j \leq n} \sum_{i=1}^n |A_{ij}|$  represents its 1-norm, and  $\|\mathbf{A}\|_{1,1} = \sum_{i,j=1}^n |A_{ij}|$  represents its  $L_{1,1}$ -norm. For a function  $h$ ,  $\nabla h$  and  $\text{grad}h$  represent its Euclidean and Riemannian gradients, respectively. Similarly,  $\nabla^2 h$  and  $\text{Hess}h$  represent its Euclidean and Riemannian Hessians, respectively. We let  $\mathbb{S}^{m-1}$  denote the unit sphere in  $\mathbb{R}^m$ . For a vector  $y$ ,  $\text{Diag}(y)$  represents the diagonal matrix whose  $i$ -th diagonal entry is  $y_i$ . Similarly for a matrix  $\mathbf{A}$ ,  $\text{diag}(\mathbf{A})$  represents the vector whose  $i$ -th entry is  $A_{ii}$ .

For brevity, we assume without loss of generality that  $\mathbf{A}$  is a symmetric matrix and  $A_{ii} = 0$ , for all  $i \in [n]$  (the latter assumption is removed in Section 4 to keep our presentation consistent with the existing works in the literature). Indeed, if  $\mathbf{A}$  is not symmetric, then we can replace  $\mathbf{A}$  by  $(\mathbf{A} + \mathbf{A}^\top)/2$ , which is a symmetric matrix, and the objective value (**Non-CVX**) remains the same for all  $\boldsymbol{\sigma} \in \mathbb{R}^{n \times r}$  since  $\boldsymbol{\sigma}\boldsymbol{\sigma}^\top$  is symmetric. Similarly, replacing the diagonal entries of  $\mathbf{A}$  by zeros decreases the objective value by the constant  $\text{Tr}(\mathbf{A})$  for all feasible  $\boldsymbol{\sigma}$ , since the diagonal entries of  $\boldsymbol{\sigma}\boldsymbol{\sigma}^\top$  are equal to 1.

The rest of the paper is organized as follows. In Section 2, we present the algorithm and discuss its per iteration cost. In Section 3, we prove the global sublinear convergence and local linear convergence of the algorithm with explicit rate estimates. In Section 4, we introduce a second-order method based on block-coordinate maximization and Lanczos method that is guaranteed to return solutions with global optimality guarantees. We also provide a global sublinear convergence rate estimate for this algorithm. We perform numerical experiments to validate our theoretical results in Section 5 and conclude the paper in Section 6.

## 2 Block-Coordinate Maximization Algorithm

In this section, we discuss the block-coordinate maximization (BCM) algorithm, its update rule and per iteration computational cost. Throughout the paper, we let  $f : \mathbb{R}^{n \times r} \rightarrow \mathbb{R}$  denote the objective function of (**Non-CVX**), i.e.,

$$f(\boldsymbol{\sigma}) = \langle \mathbf{A}, \boldsymbol{\sigma}\boldsymbol{\sigma}^\top \rangle.$$

Given the current iterate  $\boldsymbol{\sigma}^k$ , the BCM algorithm chooses a row  $i_k \in [n]$  of the matrix  $\boldsymbol{\sigma}^k$  and maximizes the following objective

$$f(\boldsymbol{\sigma}^k) = \sum_{i=1}^n \langle \sigma_i^k, g_i^k \rangle, \quad \text{where} \quad g_i^k := \sum_{j \neq i} A_{ij} \sigma_j^k,$$

over the block  $\sigma_{i_k}^k \in \mathbb{S}^{r-1}$ . More formally, we can write the update rule of the algorithm as follows

$$\begin{aligned} \sigma_{i_k}^{k+1} &= \arg \max_{\|\zeta\|=1} f(\boldsymbol{\sigma}^k) := f(\sigma_1^k, \dots, \sigma_{i_k-1}^k, \zeta, \sigma_{i_k+1}^k, \dots, \sigma_n^k), \\ &= \arg \max_{\|\zeta\|=1} 2\langle \zeta, g_{i_k}^k \rangle + \sum_{i \neq i_k} \sum_{j \neq i, i_k} A_{ij} \langle \sigma_i^k, \sigma_j^k \rangle, \\ &= \arg \max_{\|\zeta\|=1} \langle \zeta, g_{i_k}^k \rangle = \frac{g_{i_k}^k}{\|g_{i_k}^k\|}, \end{aligned} \tag{2.1}$$

---

**Algorithm 1** Block-Coordinate Maximization (BCM)

---

Initialize  $\sigma^0 \in \mathbb{R}^{n \times r}$  and calculate  $g_i^0 = \sum_{j \neq i} A_{ij} \sigma_j^0$ , for all  $i \in [n]$ .  
**for**  $k = 0, 1, 2, \dots$  **do**  
    Choose block  $i_k = i$  using one of the coordinate selection rules.  
     $\sigma_{i_k}^{k+1} \leftarrow g_{i_k}^k / \|g_{i_k}^k\|$ .  
     $g_i^{k+1} \leftarrow g_i^k - A_{ii_k} \sigma_{i_k}^k + A_{ii_k} \sigma_{i_k}^{k+1}$ , for all  $i \neq i_k$ .  
**end for**

---

with the convention that  $\sigma_{i_k}^{k+1} = \sigma_{i_k}^k$  when  $\|g_{i_k}^k\| = 0$ . Blocks  $\sigma_{i_k}^k$  that are updated at each iteration can be chosen through any deterministic or randomized rule, and in this paper we focus on three coordinate selection rules:

- Uniform sampling:  $i_k = i$  with probability  $p_i = 1/n$ .
- Importance sampling:  $i_k = i$  with probability  $p_i = \|g_i^k\| / \sum_{j=1}^n \|g_j^k\|$ .
- Greedy coordinate selection:  $i_k = \arg \max_{i \in [n]} (\|g_i^k\| - \langle \sigma_i^k, g_i^k \rangle)$ .

Per iteration computational cost of the BCM algorithm with uniform sampling is  $\mathcal{O}(nr)$  as after  $i_k$  is chosen uniformly at random,  $g_{i_k}^k$  can be computed in  $2(n-1)r$  floating point operations. On the other hand, the BCM algorithm with importance sampling and greedy coordinate selection requires all  $\{\|g_i^k\|\}_{i=1}^n$ , which can be naively computed in  $\mathcal{O}(n^2r)$  floating point operations per iteration. Instead, a smarter implementation is to keep both  $\{\sigma_i^k\}_{i=1}^n$  and  $\{g_i^k\}_{i=1}^n$ 's in the memory (only the current iterates, not all the past ones) and update them as presented in Algorithm 1, which can be done in  $2(n-1)r$  floating point operations. Therefore, per iteration computational cost of the BCM method with all three coordinate selection rules is  $\mathcal{O}(nr)$  for dense  $\mathbf{A}$  (i.e., when no structure is available on  $\mathbf{A}$ ). However, in many SDP applications (such as Max-Cut and graphical model inference),  $\mathbf{A}$  is induced by a graph and letting  $d$  denote the maximum degree of the graph that induces  $\mathbf{A}$ , the computational cost of the BCM algorithm becomes  $\mathcal{O}(dr)$ . In comparison, per iteration computational complexity of the Riemannian gradient ascent algorithm is  $\mathcal{O}(n^2r)$ , whereas the Riemannian trust-region algorithm runs a few iterations of a subroutine (e.g., power method) to solve the trust-region subproblem, whose per iteration cost is typically  $\mathcal{O}(n^2r)$ .

### 3 Convergence Rate of BCM

In this section, we analyze the convergence rate of the BCM algorithm. As the feasible set of the problem in (Non-CVX) defines a smooth manifold, we will use certain tools from manifold optimization throughout the paper, which are highlighted in the following subsection. We refer to [AMS07, Section 5.4] for a more detailed treatment of this topic.

#### 3.1 Riemannian Geometry of the Problem

We define the following submanifold of matrices  $\mathbb{R}^{n \times r}$  that corresponds to the Riemannian geometry induced by the constraints of the problem (Non-CVX) in the Euclidean space:

$$\mathcal{M}_r := \left\{ \sigma = (\sigma_1, \dots, \sigma_n)^\top \in \mathbb{R}^{n \times r} : \|\sigma_i\| = 1, \forall i \in [n] \right\}.$$

This manifold represents the Cartesian product of  $n$  unit spheres in  $\mathbb{R}^r$ . For any given point  $\boldsymbol{\sigma} \in \mathcal{M}_r$ , its tangent space can be found by taking the differential of the equality constraints as follows

$$T_{\boldsymbol{\sigma}}\mathcal{M}_r := \left\{ \mathbf{u} = (u_1, \dots, u_n)^\top \in \mathbb{R}^{n \times r} : \langle \mathbf{u}_i, \boldsymbol{\sigma}_i \rangle = 0, \forall i \in [n] \right\}.$$

The exponential map  $\text{Exp}_{\boldsymbol{\sigma}} : T_{\boldsymbol{\sigma}}\mathcal{M}_r \rightarrow \mathcal{M}_r$  on this manifold can be found as  $\boldsymbol{\sigma}' = \text{Exp}_{\boldsymbol{\sigma}}(\mathbf{u})$  such that

$$\sigma'_i = \sigma_i \cos(\|\mathbf{u}_i\|) + \frac{u_i}{\|\mathbf{u}_i\|} \sin(\|\mathbf{u}_i\|),$$

and the geodesic distance between  $\boldsymbol{\sigma}$  and  $\boldsymbol{\sigma}'$  is given by

$$\text{dist}(\boldsymbol{\sigma}, \boldsymbol{\sigma}') = \|\mathbf{u}\|_{\mathbb{F}}^2.$$

Using these definitions, the Riemannian gradient of  $f(\boldsymbol{\sigma})$  can be explicitly written as follows

$$\text{grad}f(\boldsymbol{\sigma}) = 2(\mathbf{A} - \boldsymbol{\Lambda})\boldsymbol{\sigma},$$

where  $\boldsymbol{\Lambda} = \text{Diag}(\text{diag}(\mathbf{A}\boldsymbol{\sigma}\boldsymbol{\sigma}^\top))$ . Furthermore, we have the following property for the Riemannian Hessian of  $f(\boldsymbol{\sigma})$  along the direction of a vector  $\mathbf{u} \in T_{\boldsymbol{\sigma}}\mathcal{M}_r$ :

$$\langle \mathbf{u}, \text{Hess}f(\boldsymbol{\sigma})[\mathbf{u}] \rangle = 2\langle \mathbf{u}, (\mathbf{A} - \boldsymbol{\Lambda})\mathbf{u} \rangle.$$

The detailed derivations of these can be found in Appendix A.

### 3.2 Global Rate of Convergence

In the following theorem, we consider the BCM algorithm with greedy coordinate selection and show that its functional ascent can be related to the norm of the Riemannian gradient of the function evaluated at the current iterate. By doing so, we show that the BCM algorithm returns a solution with arbitrarily small Riemannian gradient.

**Theorem 1.** *Let  $f^* = \max_{\|\sigma_i\|=1, \forall i \in [n]} f(\boldsymbol{\sigma})$ . Then, for any  $K \geq 1$ , BCM with greedy coordinate selection yields the following guarantee*

$$\min_{k \in [K-1]} \|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 \leq \frac{2n\|\mathbf{A}\|_1(f^* - f(\boldsymbol{\sigma}^0))}{K}. \quad (3.1)$$

Using a similar approach to Theorem 1, we show in the following corollary that the BCM algorithm with uniform and importance sampling attains a similar sublinear convergence rate in expectation.

**Corollary 2.** *Let  $f^* = \max_{\|\sigma_i\|=1, \forall i \in [n]} f(\boldsymbol{\sigma})$ . Then, for any  $K \geq 1$ , randomized BCM yields the following guarantee*

$$\min_{k \in [K-1]} \mathbb{E} \|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 \leq \frac{2L(f^* - f(\boldsymbol{\sigma}^0))}{K}, \quad (3.2)$$

where

$$L = \begin{cases} n\|\mathbf{A}\|_1, & \text{for uniform sampling,} \\ \|\mathbf{A}\|_{1,1}, & \text{for importance sampling.} \end{cases} \quad (3.3)$$

We can observe from (3.1) and (3.2) that BCM with uniform sampling attains the same sublinear rate as BCM with greedy coordinate selection in expectation as they both require at most  $\lceil (2n\|\mathbf{A}\|_1(f^* - f(\boldsymbol{\sigma}^0)))/\epsilon \rceil$  iterations to return a solution  $\boldsymbol{\sigma}$  satisfying  $\|\text{grad}f(\boldsymbol{\sigma})\|_{\mathbb{F}}^2 \leq \epsilon$ . On the other hand, comparing the rate guarantees in (3.2) and (3.3), we see that BCM with importance sampling enjoys a tighter convergence rate compared to BCM with uniform sampling, since  $\|\mathbf{A}\|_{1,1} \leq n\|\mathbf{A}\|_1$  for all  $\mathbf{A} \in \mathbb{R}^{n \times n}$ .



### 3.3 Local Rate of Convergence

Although the BCM algorithm enjoys the sublinear convergence rates presented in Section 3.2, it is numerically observed that the rate of convergence is linear when  $\sigma^k$  is close to a local maximum [JMRT16, WCK17]. In this section, we investigate this behavior and prove that indeed BCM attains a linear convergence rate around a local maximum under the quadratic decay condition on the objective function, which is classically defined as follows [Ani00, BI95]: Consider the unconstrained maximization problem:  $\max_x \varphi(x)$ , and let  $\Omega_{\bar{x}}$  denote the set of local maximizers with objective value  $\varphi(\bar{x})$ . Then, the quadratic decay condition is said to be satisfied at  $\bar{x}$  for  $\varphi$ , if there exists constants  $\mu, \delta > 0$  such that  $\varphi(x) \leq \varphi(\bar{x}) - \mu \text{dist}^2(x, \Omega_{\bar{x}})$ , for all  $x$  such that  $\|x - \bar{x}\| \leq \delta$ , where  $\text{dist}$  measures the distance of the point  $x$  to the set  $\Omega_{\bar{x}}$ .

For the constrained optimization problem that we are considering in (Non-CVX), this definition needs to be slightly reworked. In particular, let  $\sigma$  be a local maximum of (Non-CVX) and consider the Taylor expansion of  $\text{Exp}_\sigma(\mathbf{u})$  around  $\sigma$ :

$$f(\text{Exp}_\sigma(\mathbf{u})) = f(\sigma) + \frac{1}{2} \langle \mathbf{u}, \text{Hess}f(\sigma)[\mathbf{u}] \rangle + \mathcal{O}(\|\mathbf{u}\|_{\mathbb{F}}^3),$$

where the first-order term is zero as  $\sigma$  is a local maximum. Then, for a sufficiently small neighborhood of  $\sigma$ , the quadratic decay condition is satisfied if and only if there exists a constant  $\mu > 0$  such that  $\langle \mathbf{u}, \text{Hess}f(\sigma)[\mathbf{u}] \rangle \leq -\mu \text{dist}^2(\text{Exp}_\sigma(\mathbf{u}), \Omega_\sigma)$ , for all  $\text{Exp}_\sigma(\mathbf{u})$  sufficiently close to  $\sigma$ , where  $\Omega_\sigma$  is the set on which  $f$  has constant value  $f(\sigma)$ . Assume for the sake of simplicity that  $\sigma$  is a strict local maximum, i.e.,  $\Omega_\sigma = \{\sigma\}$ . Then, the distance between  $\text{Exp}_\sigma(\mathbf{u})$  and  $\sigma$  can be found as the norm of the tangent vector that connects these two points via the geodesic curve, i.e.,  $\text{dist}(\text{Exp}_\sigma(\mathbf{u}), \sigma) = \|\mathbf{u}\|_{\mathbb{F}}$ . Therefore, the quadratic decay condition is satisfied if and only if there exists a constant  $\mu > 0$  such that  $\langle \mathbf{u}, \text{Hess}f(\sigma)[\mathbf{u}] \rangle \leq -\mu \|\mathbf{u}\|_{\mathbb{F}}^2$  for all  $\mathbf{u} \in T_\sigma \mathcal{M}_r$ , where we note that the condition that  $\text{Exp}_\sigma(\mathbf{u})$  is sufficiently close to  $\sigma$  is dropped considering the limit as  $\mathbf{u} \rightarrow \mathbf{0}$ .

Unfortunately, no local maximum is a strict local maximum for the problem (Non-CVX). To observe this, let  $O(r) = \{\mathbf{Q} \in \mathbb{R}^{r \times r} : \mathbf{Q}^\top \mathbf{Q} = \mathbf{Q} \mathbf{Q}^\top = \mathbf{I}\}$  denote the orthogonal group in dimension  $r$ . Then, it can be observed that  $f(\sigma \mathbf{Q}) = \langle \mathbf{A}, \sigma \mathbf{Q} \mathbf{Q}^\top \sigma^\top \rangle = \langle \mathbf{A}, \sigma \sigma^\top \rangle = f(\sigma)$ , for any  $\mathbf{Q} \in O(r)$ . Therefore, in order to measure the distance between  $\text{Exp}_\sigma(\mathbf{u})$  and  $\Omega_\sigma$ , we define the following equivalence relation  $\sim$ :

$$\sigma \sim \sigma' \iff \exists \mathbf{Q} \in O(r) : \sigma = \sigma' \mathbf{Q}. \quad (3.4)$$

This equivalence relation induces a quotient space denoted by  $\mathcal{M}_r / \sim$  and we let  $[\sigma]$  denote the equivalence class of a given matrix  $\sigma \in \mathcal{M}_r$ . According to this definition,  $f$  has constant value of  $f(\sigma)$  on the set  $[\sigma]$ , i.e.,  $\Omega_\sigma = [\sigma]$ . We let  $\mathcal{V}_\sigma \subset T_\sigma \mathcal{M}_r$  denote the tangent space to the equivalence class  $[\sigma]$ , which can be found as  $\mathcal{V}_\sigma = \{\sigma \mathbf{B} : \mathbf{B} \in \mathbb{R}^{r \times r} \text{ and } \mathbf{B}^\top = -\mathbf{B}\}$ .<sup>1</sup> Therefore,  $\text{dist}(\text{Exp}_\sigma(\mathbf{u}), [\sigma]) = \|\mathbf{u}\|_{\mathbb{F}}$  if the closest point to  $\text{Exp}_\sigma(\mathbf{u})$  in  $[\sigma]$  is  $\sigma$ , or equivalently  $\text{dist}(\text{Exp}_\sigma(\mathbf{u}), [\sigma]) = \|\mathbf{u}\|_{\mathbb{F}}$  if  $\mathbf{u} \in T_\sigma \mathcal{M}_r \setminus \mathcal{V}_\sigma$ . Consequently, we say that quadratic decay is satisfied at  $\sigma$  for  $f$  if  $\text{Hess}f(\sigma)$  is negative definite on the orthogonal complement of  $\mathcal{V}_\sigma$  in  $T_\sigma \mathcal{M}_r$ . The formal statement of this definition is as follows.

**Definition 3** (Quadratic Decay). *Let  $\sigma$  be a local maximum of (Non-CVX). Quadratic decay condition is said to be satisfied at  $\sigma$  for  $f$  if there exists a constant  $\mu > 0$  such that*

$$\langle \mathbf{u}, \text{Hess}f(\sigma)[\mathbf{u}] \rangle \leq -\mu \|\mathbf{u}\|_{\mathbb{F}}^2, \quad \text{for all } \mathbf{u} \in T_\sigma \mathcal{M}_r \setminus \mathcal{V}_\sigma, \quad (3.5)$$

where  $\mathcal{V}_\sigma$  is the tangent space to the equivalence class  $[\sigma]$ .

<sup>1</sup>Note that the dimension of  $\mathcal{V}_\sigma$  depends on the rank of  $\sigma$ , and hence the quotient space is not a manifold.



In the following theorem, we present the linear convergence rate of the BCM algorithm under the quadratic decay condition. We defer the validity of this condition to Section 3.4 where we show that quadratic decay generically (over the set of matrices  $A$ ) holds for  $f$  when  $r$  is sufficiently large.

**Theorem 4.** *Let  $\bar{\sigma}$  be a limit point of the BCM algorithm and assume that  $\bar{\sigma}$  is a local maximum that satisfies the quadratic decay condition. If  $\sigma^0$  is sufficiently close to the equivalent class  $[\bar{\sigma}]$ , then the iterates generated by the BCM algorithm with greedy coordinate selection enjoy the following linear convergence rate*

$$f(\bar{\sigma}) - f(\sigma^{k+1}) \leq \left(1 - \frac{\mu}{4n^2\|\mathbf{A}\|_1}\right) \left(f(\bar{\sigma}) - f(\sigma^k)\right). \quad (3.6)$$

The linear convergence rate of the BCM algorithm with greedy coordinate selection in Theorem 4 can be extended for importance sampling and uniform sampling as we highlight in the following.

**Corollary 5.** *Let the conditions in Theorem 4 hold. Then, the iterates generated by the BCM algorithm enjoys the local linear convergence rate*

$$f(\bar{\sigma}) - \mathbb{E}f(\sigma^k) \leq (1 - \rho)^k \left(f(\bar{\sigma}) - f(\sigma^0)\right), \quad (3.7)$$

where  $\rho = \frac{\mu}{4n\|\mathbf{A}\|_{1,1}}$  for importance sampling and  $\rho = \frac{\mu}{4n^2\|\mathbf{A}\|_1}$  for uniform sampling.

### 3.4 Quadratic Decay Condition Holds Generically

In this section, we consider the quadratic decay condition, which is a condition on (**Non-CVX**), and relate it to a condition on the original problem in (**CVX**). In particular, we characterize sufficient conditions on (**CVX**) for quadratic decay to hold. We first provide some background on semidefinite programming (see for example [AHO97], for a more detailed treatment of this topic). Consider the SDP in (**CVX**):

$$\begin{aligned} & \text{maximize} && \langle \mathbf{A}, \mathbf{X} \rangle \\ & \text{subject to} && X_{ii} = 1, \text{ for } i \in [n], \\ & && \mathbf{X} \succeq 0, \end{aligned}$$

and its dual:

$$\begin{aligned} & \text{minimize} && \langle 1, y \rangle \\ & \text{subject to} && \mathbf{Z} = \text{Diag}(y) - \mathbf{A}, \\ & && \mathbf{Z} \succeq 0, \end{aligned}$$

where  $1$  is the vector of ones of appropriate size. Let  $\mathbf{X}^*$  and  $(y^*, \mathbf{Z}^*)$  denote the primal and dual optimal solutions, respectively, and let  $r^*$  denote the rank of  $\mathbf{X}^*$ . Then, there exists a  $\mathbf{Q} \in O(n)$  such that

$$\begin{aligned} \mathbf{X}^* &= \mathbf{Q} \text{Diag}(\lambda_1, \dots, \lambda_{r^*}, 0, \dots, 0) \mathbf{Q}^\top, \\ \mathbf{Z}^* &= \mathbf{Q} \text{Diag}(0, \dots, 0, \omega_{r^*+1}, \dots, \omega_n) \mathbf{Q}^\top. \end{aligned}$$

We say that *strict complementarity* holds if  $\lambda_i > 0$  for  $i = 1, \dots, r^*$  and  $\omega_j > 0$  for  $j = r^* + 1, \dots, n$ . Furthermore, let  $\mathbf{Q}_1 \in \mathbb{R}^{n \times r^*}$  and  $\mathbf{Q}_2 \in \mathbb{R}^{n \times (n-r^*)}$  respectively denote the first  $r^*$  columns and the last  $n - r^*$  columns of  $\mathbf{Q}$  and let  $q_i$  denote the  $i$ th row of  $\mathbf{Q}_1$ , i.e.,  $\mathbf{Q}_1 = [q_1, q_2, \dots, q_n]^\top$ . Then,

$(y^*, \mathbf{Z}^*)$  is *dual nondegenerate* if and only if  $\{q_1 q_1^\top, \dots, q_n q_n^\top\}$  spans  $\text{Sym}_{r^*}$ , i.e., the set of real symmetric  $r^* \times r^*$  matrices [AHO97, Theorem 3]. If  $\mathbf{X}^*$  and  $(y^*, \mathbf{Z}^*)$  satisfy strict complementarity and  $(y^*, \mathbf{Z}^*)$  is dual nondegenerate, then  $\mathbf{X}^*$  is the unique optimal primal solution (see [AHO97, Theorem 4]). Strict complementarity and dual nondegeneracy are known to hold generically (over the set of possible cost matrices  $\mathbf{A} \in \mathbb{R}^{n \times n}$ , i.e., they fail to hold only on a subset of measure zero of  $\mathbb{R}^{n \times n}$ ) as proven in [AHO97, Lemma 2]. Using these definitions, we show in the next theorem that strict complementarity and dual nondegeneracy are sufficient for quadratic decay to hold at the maximizer of (Non-CVX).

**Theorem 6.** *Suppose that  $\mathbf{X}^*$  and  $(y^*, \mathbf{Z}^*)$  are respectively primal and dual optimal solutions satisfying strict complementarity and dual nondegeneracy. If  $r \geq \text{rank}(\mathbf{X}^*)$ , then quadratic decay is satisfied for  $f$  at all  $\boldsymbol{\sigma}$  such that  $\boldsymbol{\sigma} \boldsymbol{\sigma}^\top = \mathbf{X}^*$ .*

This theorem states that quadratic decay holds for all global maxima of (Non-CVX) provided that the rank of the factorization is large enough so that the global maximum values of (CVX) and (Non-CVX) are equal to one another. For this case, the set of all global maxima is an equivalence class corresponding to a solution since strict complementarity and dual nondegeneracy imply that the primal solution of (CVX) is unique. On top of this, when  $r \geq \sqrt{2n}$ , it is known that (see [BVB16, Theorem 2]) any local maximum is global generically (i.e., for almost all cost matrices  $\mathbf{A}$ ). As strict complementarity and dual nondegeneracy also hold generically for (CVX), then consequently, when  $r \geq \sqrt{2n}$ , quadratic decay holds for all local maxima generically as we highlight in the following corollary.

**Corollary 7.** *If  $r \geq \sqrt{2n}$ , then quadratic decay holds for all local maxima generically.*

## 4 Approximately Achieving the Maximum Value of (CVX)

Our results in Section 3 show that the BCM algorithm converges with a sublinear rate to a first-order stationary solution and with a linear rate to a local maximum when initialized sufficiently close to it. In this section, we incorporate a second-order oracle to BCM in order to obtain an algorithm, which we refer to as BCM2, that returns an approximate second-order stationary point. More specifically, at the current iteration of the algorithm, if the norm of the gradient is large, we take a BCM step. Otherwise, we run a subroutine (e.g., Lanczos method) to find the leading eigenvector of the Hessian. The main motivation for designing such an algorithm is that the approximate second-order stationary solutions provide  $\mathcal{O}(1/r)$  approximation to (CVX). In particular, call  $\boldsymbol{\sigma}$  an  $\varepsilon$ -approximate concave point if  $\langle \mathbf{u}, \text{Hess} f(\boldsymbol{\sigma})[\mathbf{u}] \rangle \leq \varepsilon \langle \mathbf{u}, \mathbf{u} \rangle$ , for all  $\mathbf{u} \in T_{\boldsymbol{\sigma}} \mathcal{M}_r$ . Then, the following theorem provides an approximation ratio between the approximate concave points of (Non-CVX) and the maximum value of (CVX).

**Theorem 8** ([MMMMO17, Theorem 1]). *Let  $\boldsymbol{\sigma} \in \mathcal{M}_r$  be an  $\varepsilon$ -approximate concave point. Then, for any positive semidefinite  $\mathbf{A}$ , the following approximation ratio holds:*

$$f(\boldsymbol{\sigma}) \geq \left(1 - \frac{1}{r-1}\right) \text{SDP}(\mathbf{A}) - \frac{n}{2} \varepsilon, \quad (4.1)$$

where  $\text{SDP}(\mathbf{A})$  is the maximum value of (CVX).

This approximation ratio follows due to a generalization of the randomized rounding approach (most famously presented by [GW95]) applied to an  $\varepsilon$ -approximate concave point. In fact, it can be shown that it is not possible to find a better approximation ratio (in terms of the dependence on the rank of the factorization  $r$ ) for all problems  $\mathbf{A}$ . This result is highlighted in the following theorem.

---

**Algorithm 2** BCM2

---

```
1: Initialize  $\boldsymbol{\sigma}^0 \in \mathbb{R}^{n \times r}$  and calculate  $g_i^0 = \sum_{j \neq i} A_{ij} \sigma_j^0$ , for all  $i \in [n]$ .
2: for  $k = 0, 1, 2, \dots$  do
3:   Compute  $\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\text{F}}^2 = 2 \sum_{i=1}^n (\|g_i^k\|^2 - \langle \sigma_i^k, g_i^k \rangle^2)$ .
4:   if  $\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\text{F}}^2 > \varepsilon^3 / (1350 \|\mathbf{A}\|_1)$  then
5:      $i_k \leftarrow \arg \max_{i \in [n]} (\|g_i^k\| - \langle \sigma_i^k, g_i^k \rangle)$ 
6:      $\sigma_{i_k}^{k+1} \leftarrow g_{i_k}^k / \|g_{i_k}^k\|$ .
7:      $g_i^{k+1} \leftarrow g_i^k - A_{ii_k} \sigma_{i_k}^k + A_{ii_k} \sigma_{i_k}^{k+1}$ , for all  $i \neq i_k$ .
8:   else
9:     Find a direction  $\mathbf{u}^k \in T_{\boldsymbol{\sigma}^k} \mathcal{M}_r$  such that  $\langle \mathbf{u}^k, \text{Hess}f(\boldsymbol{\sigma}^k)[\mathbf{u}^k] \rangle \geq \lambda_{\max}(\text{Hess}f(\boldsymbol{\sigma}^k))/2$ ,
        $\langle \mathbf{u}^k, \text{grad}f(\boldsymbol{\sigma}^k) \rangle \geq 0$ , and  $\|\mathbf{u}^k\|_{\text{F}} = 1$ .
10:     $\sigma_i^{k+1} \leftarrow \sigma_i^k \cos(\|\mathbf{u}^k\|t) + \frac{u_i^k}{\|\mathbf{u}^k\|} \sin(\|\mathbf{u}^k\|t)$ , for all  $i \in [n]$ , where  $t = \varepsilon / (15 \|\mathbf{A}\|_1)$ .
11:     $g_i^{k+1} \leftarrow \sum_{j \neq i} A_{ij} \sigma_j^{k+1}$ , for all  $i \in [n]$ .
12:   end if
13: end for
```

---

**Theorem 9** ([BdOFV10, Theorems 1 & 3]). *Let  $\text{SDP}(\mathbf{A})$  be the maximum value of (CVX) and  $\text{SDP}_r(\mathbf{A})$  be the maximum value of (Non-CVX). Then, for all positive semidefinite matrices  $\mathbf{A}$ , the following approximation ratio holds:*

$$1 \geq \frac{\text{SDP}_r(\mathbf{A})}{\text{SDP}(\mathbf{A})} \geq \gamma(r) = \frac{2}{r} \left( \frac{\Gamma((r+1)/2)}{\Gamma(r/2)} \right)^2 = 1 - \Theta(1/r), \quad (4.2)$$

where  $\Gamma(z) = \int_0^\infty x^{z-1} e^{-x} dx$  is the Gamma function. Furthermore, under the unique games conjecture, there is no polynomial-time algorithm that approximates  $\text{SDP}_r(\mathbf{A})$  with an approximation ratio greater than  $\gamma(r) + \varepsilon$  for any  $\varepsilon > 0$ .

These results provide motivation to design algorithms with second-order guarantees to solve (Non-CVX) and for this reason, we propose the BCM2 algorithm (see Algorithm 2), which can be described as follows: When the Frobenius norm of the Riemannian gradient is at least as large as  $\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\text{F}}^2 > \varepsilon^3 / (1350 \|\mathbf{A}\|_1)$ , we use BCM to update the current solution. Otherwise, we assume that there is a second-order oracle that returns an update direction  $\mathbf{u}^k \in T_{\boldsymbol{\sigma}^k} \mathcal{M}_r$  such that  $\langle \mathbf{u}^k, \text{Hess}f(\boldsymbol{\sigma}^k)[\mathbf{u}^k] \rangle \geq \lambda_{\max}(\text{Hess}f(\boldsymbol{\sigma}^k))/2$ ,  $\langle \mathbf{u}^k, \text{grad}f(\boldsymbol{\sigma}^k) \rangle \geq 0$ , and  $\|\mathbf{u}^k\|_{\text{F}} = 1$ . Notice that finding a tangent vector  $\mathbf{u}^k$  that satisfy  $\langle \mathbf{u}^k, \text{Hess}f(\boldsymbol{\sigma}^k)[\mathbf{u}^k] \rangle \geq \lambda_{\max}(\text{Hess}f(\boldsymbol{\sigma}^k))/2$  and  $\|\mathbf{u}^k\|_{\text{F}} = 1$  is an eigenpair problem and can be solved efficiently using the Lanczos method. The condition  $\langle \mathbf{u}^k, \text{grad}f(\boldsymbol{\sigma}^k) \rangle \geq 0$ , on the other hand, can always be satisfied by switching the sign of  $\mathbf{u}^k$ . It is a straightforward exercise to explicitly construct such a vector and it can be found in [BAC16, Lemma 11].

We first analyze the convergence of Algorithm 2 in Theorem 10, where we assume that we have access to a subroutine that solves the eigenpair problem to the desired accuracy. We then implement the subroutine using the Lanczos algorithm (presented in Algorithm 3) and present its convergence in Theorem 12. In particular, we have the following theorem for the former case.

**Theorem 10.** *Consider Algorithm 2, where BCM is used at iteration  $k$  if  $\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\text{F}}^2 \geq \varepsilon^3 / (1350 \|\mathbf{A}\|_1)$  and a second-order step (see lines 9-11 of Algorithm 2) is taken otherwise. Let  $K_{\text{BCM}}$  denote the number of BCM epochs made and let  $K_{\text{H}}$  denote the number of second-order oracle iterations made*

such that  $K = nK_{\text{BCM}} + K_{\text{H}}$ . Then, as soon as

$$K_{\text{BCM}} + K_{\text{H}} = \left\lceil \frac{675n\|\mathbf{A}\|_1^2}{\varepsilon^2} \right\rceil, \quad (4.3)$$

Algorithm 2 is guaranteed to return a solution  $\boldsymbol{\sigma}^K$  that satisfies

$$f(\boldsymbol{\sigma}^K) \geq \left(1 - \frac{1}{r-1}\right) \text{SDP}(\mathbf{A}) - \frac{n}{2}\varepsilon, \quad (4.4)$$

where  $\text{SDP}(\mathbf{A})$  is the maximum value of (CVX).

In Theorem 10,  $K_{\text{BCM}} + K_{\text{H}}$  represents the total number of epochs to guarantee (4.4), whereas the iteration counter of the algorithm is given in terms of  $K = nK_{\text{BCM}} + K_{\text{H}}$ . This is due to the fact that, at each iteration of the BCM algorithm, a single row of  $\boldsymbol{\sigma}$  is updated and consequently  $n$  iterations of the BCM algorithm add up to an epoch. On the other hand, at each iteration of the second-order step, all entries of  $\boldsymbol{\sigma}$  are updated, and hence each second-order iteration is an epoch. In terms of the computational cost, an iteration of BCM requires  $\mathcal{O}(nr)$  operations and consequently an epoch of BCM requires  $\mathcal{O}(n^2r)$  operations, whereas the second-order direction of update is typically found approximately via a few iterations of the power method or the Lanczos method (see Theorem 12 for a more rigorous treatment of this statement), which require  $\mathcal{O}(n^2r)$  operations. Therefore, an epoch of Algorithm 2 typically has a computational complexity of  $\mathcal{O}(n^2r)$ . Furthermore, by Theorem 10, we observe that in at most  $\mathcal{O}(n\|\mathbf{A}\|_1^2/\varepsilon^2)$  epochs, Algorithm 2 returns a solution that achieves the optimal approximation ratio up to an accuracy of  $\mathcal{O}(n\varepsilon)$ . In particular, picking  $\varepsilon = 2\text{SDP}(\mathbf{A})/(n(r-1))$ , we obtain the following corollary.

**Corollary 11.** Consider the setup of Theorem 10 and set  $\varepsilon = 2\text{SDP}(\mathbf{A})/(n(r-1))$ . Then, as soon as

$$K = \left\lceil \frac{675n^3(r-1)^2\|\mathbf{A}\|_1^2}{4(\text{SDP}(\mathbf{A}))^2} \right\rceil, \quad (4.5)$$

Algorithm 2 is guaranteed to return a solution  $\boldsymbol{\sigma}^K$  that satisfies

$$f(\boldsymbol{\sigma}^K) \geq \left(1 - \frac{2}{r-1}\right) \text{SDP}(\mathbf{A}).$$

**Remark** In order to understand the total running time of BCM2, consider the following example. Let  $\mathbf{A}$  be the adjacency matrix of a random Erdos-Rényi graph on  $n$  nodes and  $\lfloor cn \rfloor$  edges. The size of the maximum cut in this graph normalized by the number of nodes can be bounded between  $[c/2+0.4\sqrt{c}, c/2+0.6\sqrt{c}]$  with high probability as  $n$  increases, for all sufficiently large  $c$  [GL14]. Since the maximum value of (CVX) is within 0.878 of the maximum cut [GW95], we can then conclude that  $\text{SDP}(\mathbf{A})/n = \mathcal{O}(c)$  with high probability. We can also observe that for this graph, the degree of a node approximately follows a Poisson distribution with mean  $2c$ , which can be approximated by a normal distribution with mean  $2c$  and variance  $\sqrt{2c}$ , for large  $c$  [GL14]. Then, we have  $\|\mathbf{A}\|_1 = \mathcal{O}(c \log n)$  with high probability. Therefore, for this problem, Corollary 11 states that in  $\tilde{\mathcal{O}}(nr^2)$  iterations (where tilde is used to hide the logarithmic dependences), Algorithm 2 returns a  $\mathcal{O}(1/r)$ -optimal solution with high probability. Per iteration computational cost of the algorithm is  $\mathcal{O}(nrc)$ , which results in a total running time of  $\tilde{\mathcal{O}}(n^2r^3c)$ . In comparison, Klein-Lu method (see [KL96, Lemma 4]) requires  $\tilde{\mathcal{O}}(n^2r^3c)$  running time and the matrix multiplicative weights method (see [AHK05, Theorem 3]) requires  $\tilde{\mathcal{O}}(n^2r^{3.5}/c)$  running time to return a  $1/r$ -optimal solution.

---

**Algorithm 3** Lanczos Method

---

- 1: Given  $\sigma$ , define  $H[\mathbf{u}] = \text{Hess}f(\sigma)[\mathbf{u}] + 4\|A\|_1\mathbf{u}$ . Initialize  $\mathbf{u}_1 \in T_\sigma\mathcal{M}_r$  such that  $\|\mathbf{u}_1\|_F = 1$ .  
Let  $\alpha_1 = \langle \mathbf{u}_1, H[\mathbf{u}_1] \rangle$  and  $\mathbf{r}_1 = H[\mathbf{u}_1] - \alpha_1\mathbf{u}_1$ .
  - 2: **for**  $\ell \geq 2$  **do**
  - 3:    $\beta_\ell = \|\mathbf{r}_{\ell-1}\|_F$
  - 4:    $\mathbf{u}_\ell = \mathbf{r}_{\ell-1}/\beta_\ell$  (If  $\beta_\ell = 0$ , pick  $\mathbf{u}_\ell \perp \text{span}(\mathbf{u}_1, \dots, \mathbf{u}_{\ell-1})$  arbitrarily)
  - 5:    $\alpha_\ell = \langle \mathbf{u}_\ell, H[\mathbf{u}_\ell] \rangle$
  - 6:    $\mathbf{r}_\ell = H[\mathbf{u}_\ell] - \alpha_\ell\mathbf{u}_\ell - \beta_\ell\mathbf{u}_{\ell-1}$
  - 7: **end for**
- 

In the description of Algorithm 2 (see line 9), we assumed that we have access to a vector in the tangent space of the current iterate, which satisfies certain second-order conditions. In Algorithm 3, we describe an efficient subroutine to find this desired tangent vector based on the Lanczos method. In particular, the Lanczos method returns a tridiagonal real symmetric matrix whose diagonal entries are  $\{\alpha_\ell\}_{\ell \geq 1}$  and off-diagonal entries are  $\{\beta_\ell\}_{\ell \geq 2}$ , where  $\ell$  denotes the iteration counter in Algorithm 3. The entire spectrum of such a symmetric tridiagonal matrix can be efficiently computed in almost linear time in the dimension of the matrix [CR13]. Consequently, letting  $y$  denote the leading eigenvector of this tridiagonal matrix, we can construct the desired tangent vector in Algorithm 2 as  $\mathbf{u}^k = \sum_{\ell \geq 1} y_\ell \mathbf{u}_\ell$ . It is well-known that after  $n(r-1)$  iterations, the Lanczos method constructs the leading eigenvector exactly (since order- $n(r-1)$  Krylov subspace spans the entire tangent space). Furthermore, it is also possible to analyze the performance of the Lanczos method with early termination [KW92]. Building on these ideas, we characterize the quality of the solution returned by Algorithms 2+3 in the following theorem.

**Theorem 12.** *Suppose in Algorithm 3, we initialize  $\mathbf{u}_1$  uniformly at random over  $T_\sigma\mathcal{M}_r$ . Let*

$$\ell^* = \left\lceil \left( \frac{1}{2} + 2\sqrt{\frac{\|A\|_1}{\varepsilon}} \right) \log \left( \frac{\left\lceil \frac{675n\|A\|_1^2}{\varepsilon^2} \right\rceil 1.648\sqrt{n(r-1)}}{\delta} \right) \right\rceil,$$

*and consider that Algorithm 3 is run for  $\min(\ell^*, n(r-1))$  iterations at each call from Algorithm 2. Then, after  $K$  iterations (defined as in (4.5)), Algorithm 2 returns a solution  $\sigma^K$  that satisfies*

$$f(\sigma^K) \geq \left(1 - \frac{1}{r-1}\right) \text{SDP}(A) - \frac{n}{2}\varepsilon,$$

*with probability at least  $1 - \delta$ .*

## 5 Numerical Experiments

In this section, we evaluate the empirical performance of the BCM algorithm with respect to the Riemannian gradient ascent (RGA) and Riemannian trust region (RTR) algorithms. All algorithms are implemented on Matlab and the experiments are run on a personal computer with 2.9 GHz processor and 16 GB memory. RGA and RTR algorithms are implemented using the Manopt package [BMAS14] with the default options and the algorithms are terminated when the maximum allowed time is achieved (can be read from the x-axis of our plots). We implement the BCM algorithm (see Algorithm 1) with the cyclic order  $(1, 2, \dots, n)$ . The initial iterate  $\sigma^0 \in \mathbb{R}^{n \times r}$  is the same for all algorithms and each row of  $\sigma^0$  is generated uniformly at random on  $\mathbb{S}^{r-1}$ . In all experiments, the cost matrix is generated as  $A = (G + G^\top)/n$ , where  $G_{ij} \sim \mathcal{N}(0, 1)$  for all  $i \neq j$ ,

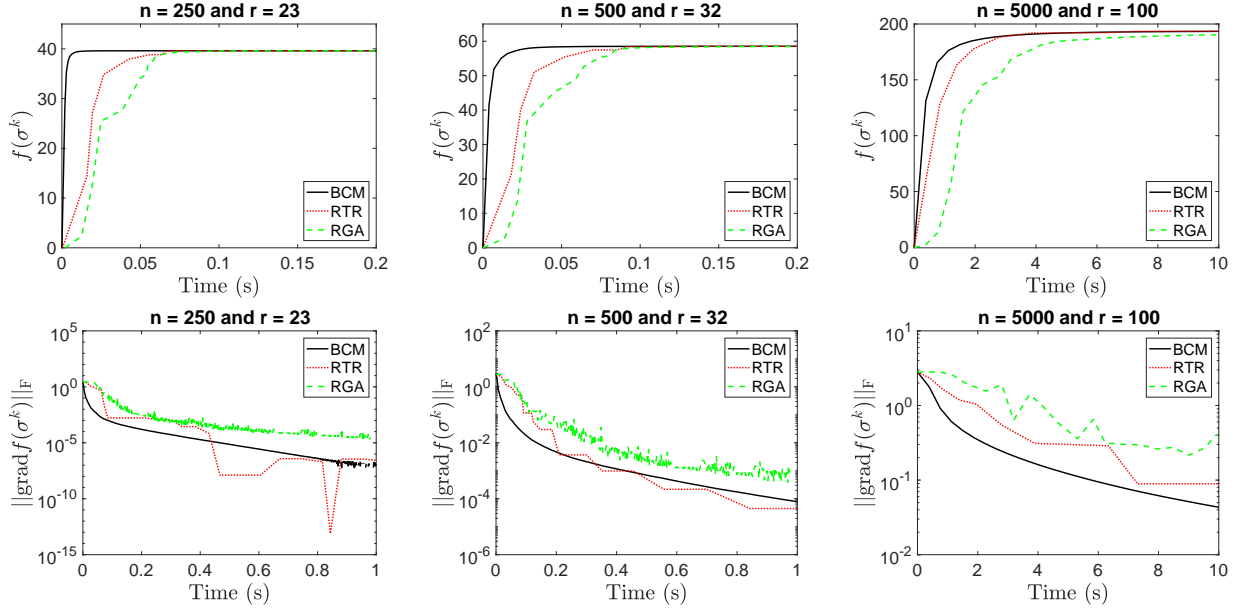


Figure 1: Objective value and gradient norm of the BCM, RGA and RTR algorithms, when  $r = \sqrt{2n}$ .

and  $G_{ii} = 0$  for all  $i \in [n]$ . We evaluate the performance of the algorithms for various values of  $n$  and  $r$ . Empirical results illustrate the fast convergence of BCM compared to RGA and RTR. We also observe that even when the problem size is large (e.g.,  $n = 20,000$ ), BCM returns a desirable solution within  $\sim 10$  seconds.

We first consider small and moderate sized problems, i.e.,  $n \in \{250, 500, 5000\}$ , and set  $r = \lceil \sqrt{2n} \rceil$  in order to recover the solution of the SDP. For the small problems, i.e., when  $n \in \{250, 500\}$ , we also solve the SDP using SDPT3 implemented on Matlab. For  $n = 250$ , the optimal value of SDP is found as 39.6147 in 4.46 seconds, whereas for  $n = 500$ , the optimal value of SDP is found as 58.5327 in 27.06 seconds. On the other hand, Burer-Monteiro factorization based algorithms are able to return the optimal solution in less than 0.1 seconds in both cases. Furthermore, we observe that for  $n = 250$ , BCM returns the optimal solution in less than 0.01 seconds, whereas RGA and RTR takes about 0.06 seconds. Similar observation can be made for  $n = 500$  case as well, and we observe that BCM is about an order of magnitude faster compared to RGA and RTR. We then observe that even the moderate-sized SDPs (e.g.,  $n = 5,000$ ) can be solved to optimality within  $\sim 10$  seconds via Burer-Monteiro approach, while CVX cannot solve the problem in 10 minutes. Among the Burer-Monteiro based methods, we observe that BCM is the fastest to converge to the optimal solution. Finally, Figure 2 shows that even when we have a large SDP of size  $n = 20,000$ , BCM returns a desirable solution in a few seconds, while RTR requires about 20 seconds and RGA requires about 50 seconds to return a solution with the same objective value.

## 6 Conclusion

In this paper, we studied the Burer-Monteiro approach to solve large-scale SDPs. We considered to solve this non-convex problem using the block-coordinate maximization algorithm that is extremely simple to implement. We proved that for various coordinate selection rules, BCM attains a global sublinear convergence rate of  $\mathcal{O}(1/\epsilon)$  to guarantee  $\mathbb{E}\|\text{grad}f(\sigma^k)\|_F^2 \leq \epsilon$ . We also showed the linear




Figure 2: Objective value and gradient norm of the BCM, RGA and RTR algorithms, when  $r \ll \sqrt{2n}$ .

convergence of BCM around a local maximum that satisfy the quadratic decay condition. We proved that the quadratic decay condition generically holds for all local maxima provided that  $r \geq \sqrt{2n}$ . These are the first precise rate estimates for the non-convex Burer-Monteiro approach in the literature to the best of our knowledge. We then introduced a new algorithm called BCM2 based on BCM and Lanczos methods. We showed that BCM2 is guaranteed to return a solution that provides  $1 - \mathcal{O}(1/r)$  approximation to the SDP without any assumptions on the cost matrix  $\mathbf{A}$ , where the  $r$ -dependence of this approximation is optimal under the unique games conjecture. We also presented numerical results that verify our theoretical findings and show that BCM is faster than the state-of-the-art methods. Even though in this paper, we only considered SDPs with diagonal constraints, it would be of interest to study the block-coordinate maximization approach in more generic problems.

## References

- [ABG07] P.-A. Absil, C.G. Baker, and K.A. Gallivan. Trust-region methods on riemannian manifolds. *Foundations of Computational Mathematics*, 7(3):303–330, Jul 2007.
- [AHK05] S. Arora, E. Hazan, and S. Kale. Fast algorithms for approximate semidefinite programming using the multiplicative weights update method. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '05, pages 339–348, 2005.
- [AHO97] F. Alizadeh, J.-P. A. Haeberly, and M. L. Overton. Complementarity and nondegeneracy in semidefinite programming. *Mathematical Programming*, 77(1):111–128, 1997.
- [AMS07] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, USA, 2007.
- [Ani00] M. Anitescu. Degenerate nonlinear programming with a quadratic growth condition. *SIAM Journal on Optimization*, 10(4):1116–1135, 2000.
- [BAC16] N. Boumal, P.-A. Absil, and C. Cartis. Global rates of convergence for nonconvex optimization on manifolds. *arXiv preprint arXiv:1605.08101*, 2016.
- [Bar95] A. I. Barvinok. Problems of distance geometry and convex properties of quadratic maps. *Discrete & Computational Geometry*, 13(2):189–202, 1995.



- [BBV16] A. S. Bandeira, N. Boumal, and V. Voroninski. On the low-rank approach for semidefinite programs arising in synchronization and community detection. *ArXiv:1602.04426*, 2016.
- [BdOFV10] J. Briët, F. M. de Oliveira Filho, and F. Vallentin. The positive semidefinite grothendieck problem with rank constraint. In *Automata, Languages and Programming*, pages 31–42, 2010.
- [BI95] J. F. Bonnans and A. Ioffe. Second-order sufficiency and quadratic growth for nonisolated minima. *Mathematics of Operations Research*, 20(4):801–817, 1995.
- [BM03] S. Burer and R. D. C. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
- [BM05] S. Burer and R. D. C. Monteiro. Local minima and convergence in low-rank semidefinite programming. *Mathematical Programming*, 103(3):427–444, Jul 2005.
- [BMAS14] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014.
- [BVB16] N. Boumal, V. Voroninski, and A. S. Bandeira. The non-convex Burer-Monteiro approach works on smooth semidefinite programs. In *Advances in Neural Information Processing Systems*, pages 2757–2765, 2016.
- [BVB18] N. Boumal, V. Voroninski, and A. S. Bandeira. Deterministic guarantees for Burer–Monteiro factorizations of smooth semidefinite programs. *arXiv preprint arXiv:1804.02008*, 2018.
- [CR13] E. S. Coakley and V. Rokhlin. A fast divide-and-conquer algorithm for computing the spectra of real symmetric tridiagonal matrices. *Applied and Computational Harmonic Analysis*, 34(3):379 – 414, 2013.
- [EDM17] M. A. Erdogdu, Y. Deshpande, and A. Montanari. Inference in graphical models via semidefinite programming hierarchies. In *Advances in Neural Information Processing Systems*, pages 416–424, 2017.
- [GH11] D. Garber and E. Hazan. Approximating semidefinite programs in sublinear time. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, pages 1080–1088, 2011.
- [GL14] D. Gamarnik and Q. Li. On the max-cut of sparse random graphs. *arXiv preprint arXiv:1411.1698*, 2014.
- [GW95] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- [JBAS10] M. Journée, F. Bach, P.-A. Absil, and R. Sepulchre. Low-rank optimization on the cone of positive semidefinite matrices. *SIAM Journal on Optimization*, 20(5):2327–2351, 2010.
- [JMRT16] A. Javanmard, A. Montanari, and F. Ricci-Tersenghi. Phase transitions in semidefinite relaxations. *Proceedings of the National Academy of Sciences*, 113(16):E2218–E2223, 2016.
- [KL96] P. Klein and H.-I Lu. Efficient approximation algorithms for semidefinite programs arising from MAX CUT and COLORING. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing, STOC ’96*, pages 338–347, New York, NY, USA, 1996. ACM.
- [KW92] J. Kuczyński and H. Woźniakowski. Estimating the largest eigenvalues by the power and lanczos algorithms with a random start. *SIAM J. Matrix Anal. Appl.*, 13(4):1094–1122, 1992.
- [LSJR16] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht. Gradient descent only converges to minimizers. In *29th Annual Conference on Learning Theory*, volume 49, pages 1246–1257. PMLR, 23–26 Jun 2016.
- [MMMO17] S. Mei, T. Misiakiewicz, A. Montanari, and R. I. Oliveira. Solving SDPs for synchronization and MaxCut problems via the Grothendieck inequality. *arXiv preprint arXiv:1703.08729*, 2017.

- [Mon16] A. Montanari. A Grothendieck-type inequality for local maxima. *arXiv preprint arXiv:1603.04064*, 2016.
- [Par03] P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96(2):293–320, May 2003.
- [Pat98] G. 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.
- [Ste10] D. Steurer. Fast SDP algorithms for constraint satisfaction problems. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 684–697, 2010.
- [TYUC17] J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher. Practical sketching algorithms for low-rank matrix approximation. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1454–1485, 2017.
- [VB96] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.
- [WCK17] P.-W. Wang, W.-C. Chang, and J. Z. Kolter. The mixing method: coordinate descent for low-rank semidefinite programming. *arXiv preprint arXiv:1706.00476*, 2017.

## A Riemannian Geometry of the Problem

Recall the definitions of the manifold and tangent space from Section 3.1:

$$\begin{aligned}\mathcal{M}_r &= \left\{ \boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)^\top \in \mathbb{R}^{n \times r} : \|\sigma_i\| = 1, \forall i \in [n] \right\}, \\ T_{\boldsymbol{\sigma}}\mathcal{M}_r &= \left\{ \mathbf{u} = (u_1, \dots, u_n)^\top \in \mathbb{R}^{n \times r} : \langle u_i, \sigma_i \rangle = 0, \forall i \in [n] \right\}.\end{aligned}$$

Before computing the Riemannian gradient and Hessian, we first let  $P^\perp : \mathbb{R}^{n \times r} \rightarrow T_{\boldsymbol{\sigma}}\mathcal{M}_r$  denote the projection operator from the Euclidean space to the tangent space of  $\boldsymbol{\sigma}$ . When applied to a given matrix  $\mathbf{w} = (w_1, \dots, w_n)^\top \in \mathbb{R}^{n \times r}$ , this projection operator yields

$$\begin{aligned}P^\perp(\mathbf{w}) &= (P_1^\perp(w_1), \dots, P_n^\perp(w_n))^\top, \\ &= (w_1 - \langle \sigma_1, w_1 \rangle \sigma_1, \dots, w_n - \langle \sigma_n, w_n \rangle \sigma_n)^\top, \\ &= \mathbf{w} - \text{Diag}(\text{diag}(\mathbf{w}\boldsymbol{\sigma}^\top)) \boldsymbol{\sigma}.\end{aligned}$$

Using this notation and standard tools from matrix manifolds [AMS07], we obtain the Riemannian gradient of  $f$  as follows

$$\text{grad}f(\boldsymbol{\sigma}) = P^\perp(\nabla f(\boldsymbol{\sigma})) = 2(\mathbf{A} - \boldsymbol{\Lambda})\boldsymbol{\sigma},$$

where  $\boldsymbol{\Lambda} = \text{Diag}(\text{diag}(\mathbf{A}\boldsymbol{\sigma}\boldsymbol{\sigma}^\top))$ . Opening up the terms in the above equality, we obtain

$$\begin{aligned}\text{grad}f(\boldsymbol{\sigma}) &= 2 \left( \mathbf{A} - \text{Diag} \left( \text{diag} \left( \mathbf{A} \begin{bmatrix} 1 & \langle \sigma_1, \sigma_2 \rangle & \dots & \langle \sigma_1, \sigma_n \rangle \\ \langle \sigma_2, \sigma_1 \rangle & 1 & \dots & \langle \sigma_2, \sigma_n \rangle \\ \vdots & \vdots & \ddots & \vdots \\ \langle \sigma_n, \sigma_1 \rangle & \langle \sigma_n, \sigma_2 \rangle & \dots & 1 \end{bmatrix} \right) \right) \right) \boldsymbol{\sigma}, \\ &= 2 \left( \mathbf{A} - \begin{bmatrix} \langle \sigma_1, g_1 \rangle & & & \\ & \ddots & & \\ & & \langle \sigma_n, g_n \rangle & \end{bmatrix} \right) \boldsymbol{\sigma}, \\ &= 2 \begin{bmatrix} -\langle \sigma_1, g_1 \rangle & A_{12} & \dots & A_{1n} \\ A_{21} & -\langle \sigma_2, g_2 \rangle & \dots & A_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ A_{n1} & A_{n2} & \dots & -\langle \sigma_n, g_n \rangle \end{bmatrix} \boldsymbol{\sigma}.\end{aligned}$$

Hence, the Riemannian gradient can be explicitly written as follows

$$\text{grad}f(\boldsymbol{\sigma}) = 2(g_1 - \langle \sigma_1, g_1 \rangle \sigma_1, \dots, g_n - \langle \sigma_n, g_n \rangle \sigma_n)^\top.$$

In particular, we have

$$\|\text{grad}f(\boldsymbol{\sigma})\|_{\mathbb{F}}^2 = 2 \sum_{i=1}^n \|g_i - \langle \sigma_i, g_i \rangle \sigma_i\|^2 = 2 \sum_{i=1}^n (\|g_i\|^2 - \langle \sigma_i, g_i \rangle^2). \quad (\text{A.1})$$

Using the same approach, we can calculate the Riemannian Hessian of  $f(\boldsymbol{\sigma})$  along the direction of a vector  $\mathbf{u} \in T_{\boldsymbol{\sigma}}\mathcal{M}_r$  by projecting the directional derivative of the gradient vector field onto the tangent space of  $\boldsymbol{\sigma}$  as follows

$$\text{Hess}f(\boldsymbol{\sigma})[\mathbf{u}] = P^\perp(\text{D grad}f(\boldsymbol{\sigma})[\mathbf{u}]),$$

where  $D \text{grad} f(\boldsymbol{\sigma})[\mathbf{u}]$  denotes the directional gradient of  $\text{grad} f(\boldsymbol{\sigma})$  along the direction  $\mathbf{u}$ . This yields

$$\text{Hess} f(\boldsymbol{\sigma})[\mathbf{u}] = \mathbf{P}^\perp \left( 2(\mathbf{A} - \boldsymbol{\Lambda})\mathbf{u} - 2 \text{ddiag}(\mathbf{A}\boldsymbol{\sigma}\mathbf{u}^\top + \mathbf{A}\mathbf{u}\boldsymbol{\sigma}^\top) \boldsymbol{\sigma} \right) = \mathbf{P}^\perp(2(\mathbf{A} - \boldsymbol{\Lambda})\mathbf{u}). \quad (\text{A.2})$$

In particular, we have

$$\langle \mathbf{u}, \text{Hess} f(\boldsymbol{\sigma})[\mathbf{u}] \rangle = 2\langle \mathbf{u}, (\mathbf{A} - \boldsymbol{\Lambda})\mathbf{u} \rangle, \quad (\text{A.3})$$

for any  $\mathbf{u} \in T_{\boldsymbol{\sigma}}\mathcal{M}_r$ .

The geodesics  $t \mapsto \boldsymbol{\sigma}(t)$  (i.e., curves of shortest path with zero acceleration) can be expressed as a function of  $\boldsymbol{\sigma} = \boldsymbol{\sigma}(0) \in \mathcal{M}_r$  and  $\mathbf{u} \in T_{\boldsymbol{\sigma}}\mathcal{M}_r$  as follows

$$\sigma_i(t) = \sigma_i \cos(\|u_i\|t) + \frac{u_i}{\|u_i\|} \sin(\|u_i\|t). \quad (\text{A.4})$$

This geodesic can be thought as the curve on the manifold that are obtained by moving from  $\boldsymbol{\sigma} \in \mathcal{M}_r$  towards the direction pointed by  $\mathbf{u} \in T_{\boldsymbol{\sigma}}\mathcal{M}_r$ . According to this definition, the exponential map  $\text{Exp}_{\boldsymbol{\sigma}} : T_{\boldsymbol{\sigma}}\mathcal{M}_r \rightarrow \mathcal{M}_r$  corresponds to evaluating the point at  $t = 1$  on the geodesic function, i.e., letting  $\boldsymbol{\sigma}' = \text{Exp}_{\boldsymbol{\sigma}}(\mathbf{u})$ , where  $\mathbf{u} \in T_{\boldsymbol{\sigma}}\mathcal{M}_r$ , we have

$$\sigma'_i = \sigma_i \cos(\|u_i\|) + \frac{u_i}{\|u_i\|} \sin(\|u_i\|).$$

These geodesics also yield the following distance function defined on the manifold

$$\text{dist}(\boldsymbol{\sigma}, \boldsymbol{\sigma}') = \left( \sum_{i=1}^n (\arccos \langle \sigma_i, \sigma'_i \rangle)^2 \right)^{1/2}, \quad (\text{A.5})$$

where letting  $\boldsymbol{\sigma}' = \text{Exp}_{\boldsymbol{\sigma}}(\mathbf{u})$ , we obtain

$$\begin{aligned} \text{dist}(\boldsymbol{\sigma}, \boldsymbol{\sigma}') &= \left( \sum_{i=1}^n (\arccos \langle \sigma_i, \sigma_i \cos \|u_i\| \rangle)^2 \right)^{1/2}, \\ &= \left( \sum_{i=1}^n \|u_i\|^2 \right)^{1/2}, \\ &= \|\mathbf{u}\|_{\text{F}}. \end{aligned}$$

Similarly, the distance between a point  $\boldsymbol{\sigma}$  and an equivalence class  $[\boldsymbol{\sigma}']$  can be found as

$$\text{dist}(\boldsymbol{\sigma}, [\boldsymbol{\sigma}']) = \min_{\mathbf{Q} \in O(r)} \text{dist}(\boldsymbol{\sigma}, \boldsymbol{\sigma}'\mathbf{Q}).$$

## B An Ascent Lemma for BCM

**Lemma 13.** *Suppose at the  $k$ -th iteration of the BCM algorithm,  $i_k$ -th block is chosen (with some coordinate selection rule). Then, the BCM algorithm yields the following ascent on the objective value:*

$$f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) = 2 \left( \|g_{i_k}^k\| - \langle \sigma_{i_k}^k, g_{i_k}^k \rangle \right) \geq 0.$$

**Proof** According to the update rule of the BCM algorithm, we have  $\sigma_{i_k}^{k+1} = \frac{g_{i_k}^k}{\|g_{i_k}^k\|}$ , which yields

$$\begin{aligned}
f(\boldsymbol{\sigma}^{k+1}) &= \sum_{i=1}^n \langle \sigma_i^{k+1}, g_i^{k+1} \rangle, \\
&= \langle \sigma_{i_k}^{k+1}, g_{i_k}^{k+1} \rangle + \sum_{i \neq i_k} \langle \sigma_i^{k+1}, g_i^{k+1} \rangle, \\
&= \|g_{i_k}^k\| + \sum_{i \neq i_k} \langle \sigma_i^k, g_i^k - A_{ii_k} \sigma_{i_k}^k + A_{ii_k} \frac{g_{i_k}^k}{\|g_{i_k}^k\|} \rangle, \tag{B.1}
\end{aligned}$$

where the last equality follows since  $g_{i_k}^{k+1} = g_{i_k}^k$  as this quantity is independent of  $\sigma_{i_k}^{k+1}$  and we have the following for all  $i \neq i_k$ :

$$\begin{aligned}
g_i^{k+1} &= \sum_{j \neq i} A_{ij} \sigma_j^{k+1}, \\
&= A_{ii_k} \sigma_{i_k}^{k+1} + \sum_{j \neq i, i_k} A_{ij} \sigma_j^{k+1}, \\
&= A_{ii_k} \frac{g_{i_k}^k}{\|g_{i_k}^k\|} + \sum_{j \neq i, i_k} A_{ij} \sigma_j^k, \\
&= A_{ii_k} \frac{g_{i_k}^k}{\|g_{i_k}^k\|} + g_i^k - A_{ii_k} \sigma_{i_k}^k.
\end{aligned}$$

Separating the terms in the sum in (B.1) and using the fact that  $A$  is a symmetric matrix, we obtain

$$\begin{aligned}
f(\boldsymbol{\sigma}^{k+1}) &= \|g_{i_k}^k\| + \sum_{i \neq i_k} A_{ii_k} \langle \sigma_i^k, \frac{g_{i_k}^k}{\|g_{i_k}^k\|} - \sigma_{i_k}^k \rangle + \sum_{i \neq i_k} \langle \sigma_i^k, g_i^k \rangle, \\
&= \|g_{i_k}^k\| + \langle g_{i_k}^k, \frac{g_{i_k}^k}{\|g_{i_k}^k\|} - \sigma_{i_k}^k \rangle + \sum_{i \neq i_k} \langle \sigma_i^k, g_i^k \rangle, \\
&= 2\|g_{i_k}^k\| - \langle \sigma_{i_k}^k, g_{i_k}^k \rangle + \sum_{i \neq i_k} \langle \sigma_i^k, g_i^k \rangle, \\
&= f(\boldsymbol{\sigma}^k) + 2\left(\|g_{i_k}^k\| - \langle \sigma_{i_k}^k, g_{i_k}^k \rangle\right),
\end{aligned}$$

which concludes the proof of the lemma.  $\square$

## C Proof of Theorem 1

From Lemma 13, we have

$$f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) = 2\left(\|g_{i_k}^k\| - \langle \sigma_{i_k}^k, g_{i_k}^k \rangle\right) = 2 \max_{i \in [n]} \left(\|g_i^k\| - \langle \sigma_i^k, g_i^k \rangle\right),$$

where the last equality follows by the greedy coordinate selection rule. We can rewrite this equality as follows:

$$\begin{aligned} f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) &= \max_{i \in [n]} \frac{2\|g_i^k\| \left( \|g_i^k\| - \langle \sigma_i^k, g_i^k \rangle \right)}{\|g_i^k\|}, \\ &\geq \max_{i \in [n]} \frac{\|g_i^k\|^2 - \langle \sigma_i^k, g_i^k \rangle^2}{\|g_i^k\|}, \end{aligned}$$

where the inequality follows since  $\|g_i^k\| \geq \langle \sigma_i^k, g_i^k \rangle$ , for all  $\sigma_i^k \in \mathbb{R}^{n \times r}$ . Lower bounding the maximum with the mean of the terms, we get

$$f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) \geq \sum_{i=1}^n \frac{\|g_i^k\|^2 - \langle \sigma_i^k, g_i^k \rangle^2}{n\|g_i^k\|}. \quad (\text{C.1})$$

The  $\|g_i^k\|$  term in the denominator in (C.1) can be upper bounded as follows

$$\|g_{i_k}^k\| \leq \sum_{j \neq i_k} |A_{i_k j}| \|\sigma_j^k\| \leq \|\mathbf{A}\|_1. \quad (\text{C.2})$$

Using this bound in (C.1), we get

$$f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) \geq \frac{1}{n\|\mathbf{A}\|_1} \sum_{i=1}^n \left( \|g_i^k\|^2 - \langle \sigma_i^k, g_i^k \rangle^2 \right) = \frac{\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2}{2n\|\mathbf{A}\|_1}. \quad (\text{C.3})$$

In order to prove (3.1), we assume the contrary that  $\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 > \epsilon$ , for all  $k \in [K-1]$ . Then, using the boundedness of  $f$ , we observe that

$$f^* - f(\boldsymbol{\sigma}^0) \geq f(\boldsymbol{\sigma}^K) - f(\boldsymbol{\sigma}^0) = \sum_{k=0}^{K-1} \left[ f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) \right].$$

Using the functional ascent of BCM in (C.3) above, we get

$$f^* - f(\boldsymbol{\sigma}^0) \geq \sum_{k=0}^{K-1} \frac{\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2}{2n\|\mathbf{A}\|_1} > \frac{K\epsilon}{2n\|\mathbf{A}\|_1},$$

where the last inequality follows by the assumption. Then, by contradiction, the algorithm returns a solution with  $\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 \leq \epsilon$ , for some  $k \in [K-1]$ , provided that

$$K \geq \frac{2n\|\mathbf{A}\|_1(f^* - f(\boldsymbol{\sigma}^0))}{\epsilon}.$$

## D Proof of Corollary 2

Similar to the proof of Theorem 1, from Proposition 13, we have

$$\begin{aligned} f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) &= 2 \left( \|g_{i_k}^k\| - \langle \sigma_{i_k}^k, g_{i_k}^k \rangle \right), \\ &= \frac{2\|g_{i_k}^k\| \left( \|g_{i_k}^k\| - \langle \sigma_{i_k}^k, g_{i_k}^k \rangle \right)}{\|g_{i_k}^k\|}, \\ &\geq \frac{\|g_{i_k}^k\|^2 - \langle \sigma_{i_k}^k, g_{i_k}^k \rangle^2}{\|g_{i_k}^k\|}, \end{aligned} \quad (\text{D.1})$$

where the inequality follows since  $\|g_{i_k}^k\| \geq \langle \sigma_{i_k}^k, g_{i_k}^k \rangle$ , for all  $\sigma_{i_k}^k \in \mathbb{R}^{n \times r}$ . Letting  $\mathbb{E}_k$  denote the expectation over  $i_k$  given  $\sigma^k$ , we have

$$\mathbb{E}_k f(\sigma^{k+1}) - f(\sigma^k) \geq \sum_{i=1}^n p_i \frac{\|g_i^k\|^2 - \langle \sigma_i^k, g_i^k \rangle^2}{\|g_i^k\|}.$$

In particular, when  $p_i = \frac{1}{n}$ , for all  $i \in [n]$  (i.e., for uniform sampling case), we have

$$\mathbb{E}_k f(\sigma^{k+1}) - f(\sigma^k) \geq \frac{1}{n\|\mathbf{A}\|_1} \sum_{i=1}^n (\|g_i^k\|^2 - \langle \sigma_i^k, g_i^k \rangle^2),$$

since  $\|g_i^k\| \leq \|\mathbf{A}\|_1$ , for all  $i \in [n]$  by (C.2). Therefore, we have

$$\mathbb{E}_k f(\sigma^{k+1}) - f(\sigma^k) \geq \frac{\|\text{grad}f(\sigma^k)\|_{\mathbb{F}}^2}{2n\|\mathbf{A}\|_1}. \quad (\text{D.2})$$

On the other hand, when  $p_i = \frac{\|g_i^k\|}{\sum_{j=1}^n \|g_j^k\|}$  (i.e., for importance sampling case), we have

$$\mathbb{E}_k f(\sigma^{k+1}) - f(\sigma^k) \geq \frac{\sum_{i=1}^n \|g_i^k\|^2 - \langle \sigma_i^k, g_i^k \rangle^2}{\sum_{j=1}^n \|g_j^k\|} = \frac{\|\text{grad}f(\sigma^k)\|_{\mathbb{F}}^2}{2\sum_{j=1}^n \|g_j^k\|}.$$

Letting  $\|\mathbf{A}\|_{1,1} = \sum_{i,j=1}^n |\mathbf{A}_{ij}|$  denote the  $L_{1,1}$  norm of matrix  $\mathbf{A}$ , we observe that  $\sum_{j=1}^n \|g_j^k\| \leq \|\mathbf{A}\|_{1,1}$ , which in the above inequality yields

$$\mathbb{E}_k f(\sigma^{k+1}) - f(\sigma^k) \geq \frac{\|\text{grad}f(\sigma^k)\|_{\mathbb{F}}^2}{2\|\mathbf{A}\|_{1,1}}. \quad (\text{D.3})$$

In order to prove (3.2), which corresponds to uniform sampling case, we assume the contrary that  $\mathbb{E}\|\text{grad}f(\sigma^k)\|_{\mathbb{F}}^2 > \epsilon$ , for all  $k \in [K-1]$ . Then, using the boundedness of  $f$ , we get

$$f^* - f(\sigma^0) \geq \mathbb{E}f(\sigma^K) - f(\sigma^0) = \sum_{k=0}^{K-1} \mathbb{E}[f(\sigma^{k+1}) - f(\sigma^k)] = \sum_{k=0}^{K-1} \mathbb{E}[\mathbb{E}_k f(\sigma^{k+1}) - f(\sigma^k)].$$

Using the expected functional ascent of BCM in (D.2) above, we get

$$f^* - f(\sigma^0) \geq \sum_{k=0}^{K-1} \frac{\mathbb{E}\|\text{grad}f(\sigma^k)\|_{\mathbb{F}}^2}{2n\|\mathbf{A}\|_1} > \frac{K\epsilon}{2n\|\mathbf{A}\|_1}, \quad (\text{D.4})$$

where the last inequality follows by the assumption. Then, by contradiction, the algorithm returns a solution with  $\mathbb{E}\|\text{grad}f(\sigma^k)\|_{\mathbb{F}}^2 \leq \epsilon$ , for some  $k \in [K-1]$ , provided that

$$K \geq \frac{2n\|\mathbf{A}\|_1(f^* - f(\sigma^0))}{\epsilon}.$$

The proof of (3.3), which corresponds to importance sampling case, can be obtained by using (D.3) (instead of (D.2)) in (D.4), and hence is omitted.



## E Proof of Theorem 4

By (C.3), we have the following functional ascent

$$f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) \geq \frac{\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2}{2n\|\mathbf{A}\|_1}. \quad (\text{E.1})$$

In order to prove linear convergence, our aim is to show that  $\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 \geq c(f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k))$ , for some  $0 < c < 2n\|\mathbf{A}\|_1$ , in a neighborhood around the limit points of the iterates generated by the algorithm. Let  $\bar{\boldsymbol{\sigma}}$  be the limit point of a subsequence  $\{\boldsymbol{\sigma}^{k_\ell}\}_{k_\ell \geq 0}$  that contains  $\boldsymbol{\sigma}^k$ . Then, we consider the solution  $\boldsymbol{\sigma} \in [\bar{\boldsymbol{\sigma}}]$  such that  $\boldsymbol{\sigma}$  is the projection of  $\boldsymbol{\sigma}^k$  onto  $[\bar{\boldsymbol{\sigma}}]$ , i.e.,  $d(\boldsymbol{\sigma}, \boldsymbol{\sigma}^k) \leq d(\boldsymbol{\sigma}', \boldsymbol{\sigma}^k)$ , for all  $\boldsymbol{\sigma}' \in [\bar{\boldsymbol{\sigma}}]$ . Then, by construction there exists  $\bar{\mathbf{u}} \perp \mathcal{V}_{\boldsymbol{\sigma}}$  such that  $\text{Exp}_{\boldsymbol{\sigma}}(\bar{\mathbf{u}}) = \boldsymbol{\sigma}^k$ . In order to perform the local convergence analysis in a unified manner, we let  $\mathbf{u} = \bar{\mathbf{u}}/\|\bar{\mathbf{u}}\|_{\mathbb{F}}$  denote the normalized tangent vector and consider the following geodesic to describe  $\boldsymbol{\sigma}^k$ :

$$\sigma_i^k = \sigma_i \cos(\|u_i\|t) + \frac{u_i}{\|u_i\|} \sin(\|u_i\|t), \quad (\text{E.2})$$

where it can be observed that  $t = \|\bar{\mathbf{u}}\|_{\mathbb{F}}$  recovers the original exponential map  $\boldsymbol{\sigma}^k = \text{Exp}_{\boldsymbol{\sigma}}(\bar{\mathbf{u}})$ . The second order Taylor approximation to (E.2) yields (note that  $t = \|\bar{\mathbf{u}}\|_{\mathbb{F}} < 1$ , when  $\boldsymbol{\sigma}$  and  $\boldsymbol{\sigma}^k$  are sufficiently close):

$$\sigma_i^k = \sigma_i + tu_i - \frac{t^2}{2}\|u_i\|^2\sigma_i + O(t^3),$$

and using this approximation, we obtain

$$g_i^k = g_i + tv_i - \frac{t^2}{2}\tilde{g}_i + O(t^3),$$

where

$$v_i^k = \sum_{j \neq i} A_{ij}u_j \quad \text{and} \quad \tilde{g}_i = \sum_{j \neq i} A_{ij}\|u_j\|^2\sigma_j.$$

This yields the following Taylor approximation to  $\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2$ :

$$\begin{aligned} \|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 &= 2 \sum_{i=1}^n \left( \|g_i^k\|^2 - \langle \sigma_i^k, g_i^k \rangle^2 \right) \\ &= 2 \sum_{i=1}^n \left( \|g_i + tv_i - \frac{t^2}{2}\tilde{g}_i\|^2 - \langle \sigma_i + tu_i - \frac{t^2}{2}\|u_i\|^2\sigma_i, g_i + tv_i - \frac{t^2}{2}\tilde{g}_i \rangle^2 \right) + O(t^3), \\ &= 2 \sum_{i=1}^n \left\{ \|g_i\|^2 + 2t\langle g_i, v_i \rangle - t^2\langle g_i, \tilde{g}_i \rangle + t^2\|v_i\|^2 \right. \\ &\quad \left. - \left( \langle \sigma_i, g_i \rangle + t\langle \sigma_i, v_i \rangle - \frac{t^2}{2}\langle \sigma_i, \tilde{g}_i \rangle + t\langle u_i, g_i \rangle + t^2\langle u_i, v_i \rangle - \frac{t^2}{2}\|u_i\|^2\langle \sigma_i, g_i \rangle \right)^2 \right\} + O(t^3). \end{aligned}$$

Note that we have  $\sigma_i = g_i/\|g_i\|$  due to the property of fixed points of the BCM algorithm and also  $\langle \sigma_i, u_i \rangle = 0$  as  $u \in T_{\sigma} \mathcal{M}_r$ . Using these relations in the above equality, we get

$$\begin{aligned}
\|\text{grad} f(\boldsymbol{\sigma}^k)\|_{\text{F}}^2 &= 2 \sum_{i=1}^n \left\{ \|g_i\|^2 + 2t\|g_i\|\langle \sigma_i, v_i \rangle - t^2\|g_i\|\langle \sigma_i, \tilde{g}_i \rangle + t^2\|v_i\|^2 \right. \\
&\quad \left. - \left( \|g_i\| + t\langle \sigma_i, v_i \rangle - \frac{t^2}{2}\langle \sigma_i, \tilde{g}_i \rangle + t^2\langle u_i, v_i \rangle - \frac{t^2}{2}\|u_i\|^2\|g_i\| \right)^2 \right\} + O(t^3), \\
&= 2 \sum_{i=1}^n \left\{ \|g_i\|^2 + 2t\|g_i\|\langle \sigma_i, v_i \rangle - t^2\|g_i\|\langle \sigma_i, \tilde{g}_i \rangle + t^2\|v_i\|^2 \right. \\
&\quad \left. - \left( \|g_i\|^2 + 2t\|g_i\|\langle \sigma_i, v_i \rangle - t^2\|g_i\|\langle \sigma_i, \tilde{g}_i \rangle + 2t^2\|g_i\|\langle u_i, v_i \rangle \right. \right. \\
&\quad \quad \left. \left. - t^2\|u_i\|^2\|g_i\|^2 + t^2\langle \sigma_i, v_i \rangle^2 \right) \right\} + O(t^3), \\
&= 2t^2 \sum_{i=1}^n \left( \|v_i\|^2 - \langle \sigma_i, v_i \rangle^2 - 2\|g_i\|\langle u_i, v_i \rangle + \|u_i\|^2\|g_i\|^2 \right) + O(t^3). \tag{E.3}
\end{aligned}$$

Since  $\langle \sigma_i, u_i \rangle = 0$ , we have by the Pythagorean theorem that

$$\|v_i\|^2 - \langle \sigma_i, v_i \rangle^2 - \left\langle \frac{u_i}{\|u_i\|}, v_i \right\rangle^2 \geq 0.$$

Using this inequality in (E.3), we get

$$\begin{aligned}
\|\text{grad} f(\boldsymbol{\sigma}^k)\|_{\text{F}}^2 &\geq 2t^2 \sum_{i=1}^n \left( \left\langle \frac{u_i}{\|u_i\|}, v_i \right\rangle^2 - 2\|g_i\|\langle u_i, v_i \rangle + \|u_i\|^2\|g_i\|^2 \right) + O(t^3), \\
&= 2t^2 \sum_{i=1}^n \left( \|u_i\|\|g_i\| - \left\langle \frac{u_i}{\|u_i\|}, v_i \right\rangle \right)^2 + O(t^3). \tag{E.4}
\end{aligned}$$

In order to lower bound (E.4) by  $c(f(\boldsymbol{\sigma}) - f(\boldsymbol{\sigma}^k))$ , we consider the second order Taylor approximation of  $f(\boldsymbol{\sigma}^k)$ , which can be written as follows

$$\begin{aligned}
f(\boldsymbol{\sigma}^k) &= \sum_{i=1}^n \langle \sigma_i^k, g_i^k \rangle, \\
&= \sum_{i=1}^n \langle \sigma_i + tu_i - \frac{t^2}{2}\|u_i\|^2\sigma_i, g_i + tv_i - \frac{t^2}{2}\tilde{g}_i \rangle + O(t^3), \\
&= \sum_{i=1}^n \left( \langle \sigma_i, g_i \rangle + t\langle \sigma_i, v_i \rangle - \frac{t^2}{2}\langle \sigma_i, \tilde{g}_i \rangle + t\langle u_i, g_i \rangle + t^2\langle u_i, v_i \rangle - \frac{t^2}{2}\|u_i\|^2\langle \sigma_i, g_i \rangle \right) + O(t^3).
\end{aligned}$$

Similar to the previous derivations, using the fact that  $\sigma_i = \frac{g_i}{\|g_i\|}$  and  $\langle \sigma_i, u_i \rangle = 0$ , for all  $i \in [n]$ , we obtain

$$\begin{aligned}
f(\boldsymbol{\sigma}^k) &= f(\boldsymbol{\sigma}) + \sum_{i=1}^n \left( t \langle \sigma_i, v_i \rangle - \frac{t^2}{2} \langle \sigma_i, \tilde{g}_i \rangle + t^2 \langle u_i, v_i \rangle - \frac{t^2}{2} \|u_i\|^2 \langle \sigma_i, g_i \rangle \right) + O(t^3), \\
&= f(\boldsymbol{\sigma}) + \sum_{i=1}^n \left( t \sum_{j \neq i} A_{ij} \langle \sigma_i, u_j \rangle - \frac{t^2}{2} \sum_{j \neq i} A_{ij} \|u_j\|^2 \langle \sigma_i, \sigma_j \rangle + t^2 \langle u_i, v_i \rangle - \frac{t^2}{2} \|u_i\|^2 \langle \sigma_i, g_i \rangle \right) + O(t^3), \\
&= f(\boldsymbol{\sigma}) + t \sum_{j=1}^n \sum_{i \neq j} A_{ji} \langle \sigma_i, u_j \rangle - \frac{t^2}{2} \sum_{j=1}^n \sum_{i \neq j} A_{ji} \|u_j\|^2 \langle \sigma_i, \sigma_j \rangle \\
&\quad + t^2 \sum_{i=1}^n \left( \langle u_i, v_i \rangle - \frac{1}{2} \|u_i\|^2 \langle \sigma_i, g_i \rangle \right) + O(t^3),
\end{aligned}$$

where the last line follows since  $\mathbf{A}$  is symmetric. Using the definition  $g_j = \sum_{i \neq j} A_{ji} \sigma_i$  and  $\sigma_i = \frac{g_i}{\|g_i\|}$  in the above inequality yields

$$\begin{aligned}
f(\boldsymbol{\sigma}^k) &= f(\boldsymbol{\sigma}) + t \sum_{j=1}^n \langle g_j, u_j \rangle - \frac{t^2}{2} \sum_{j=1}^n \|u_j\|^2 \langle g_j, \sigma_j \rangle + t^2 \sum_{i=1}^n \left( \langle u_i, v_i \rangle - \frac{1}{2} \|u_i\|^2 \langle \sigma_i, g_i \rangle \right) + O(t^3), \\
&= f(\boldsymbol{\sigma}) + t^2 \sum_{i=1}^n \left( \langle u_i, v_i \rangle - \|u_i\|^2 \|g_i\| \right) + O(t^3). \tag{E.5}
\end{aligned}$$

Reorganizing terms, we get

$$f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k) = f(\boldsymbol{\sigma}) - f(\boldsymbol{\sigma}^k) = t^2 \sum_{i=1}^n \left( \|u_i\|^2 \|g_i\| - \langle u_i, v_i \rangle \right) + O(t^3). \tag{E.6}$$

Turning back our attention to (E.4), we can lower bound the right-hand side as follows

$$\begin{aligned}
\|\text{grad} f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 &\geq 2t^2 \sum_{i=1}^n \frac{1}{\|u_i\|^2} \left( \|u_i\|^2 \|g_i\| - \langle u_i, v_i \rangle \right)^2 + O(t^3), \\
&\geq 2t^2 \sum_{i=1}^n \left( \|u_i\|^2 \|g_i\| - \langle u_i, v_i \rangle \right)^2 + O(t^3), \\
&\geq \frac{2t^2}{n} \left( \sum_{i=1}^n \left( \|u_i\|^2 \|g_i\| - \langle u_i, v_i \rangle \right) \right)^2 + O(t^3),
\end{aligned}$$

where the second inequality follows since  $\|u_i\|^2 \leq \|\mathbf{u}\|_{\mathbb{F}}^2 = 1$  and the last inequality follows since  $(\sum_{i=1}^n a_i)^2 \leq n \sum_{i=1}^n a_i^2$ , for all  $a_i \in \mathbb{R}$ ,  $i \in [n]$ . Using the second order approximation derived in (E.6) in the above inequality, we obtain

$$\begin{aligned}
\|\text{grad} f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 &\geq \frac{f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k)}{n} \sum_{i=1}^n 2 \left( \|u_i\|^2 \|g_i\| - \langle u_i, v_i \rangle \right) + O(t^3), \\
&= \frac{2 \langle \mathbf{u}, (\mathbf{\Lambda} - \mathbf{A}) \mathbf{u} \rangle}{n} \left( f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k) \right) + O(t^3),
\end{aligned}$$

where  $\mathbf{\Lambda} = \text{Diag}(\|g_1\|, \dots, \|g_n\|)$ . Since we have  $2 \langle \mathbf{u}, (\mathbf{A} - \mathbf{\Lambda}) \mathbf{u} \rangle \leq -\mu \|\mathbf{u}\|_{\mathbb{F}}^2$  by the quadratic decay condition, we conclude that

$$\|\text{grad} f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 \geq \frac{\mu}{n} \left( f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k) \right) + O(t^3). \tag{E.7}$$

This implies that whenever  $\sigma^k$  is sufficiently close to  $\sigma$ , i.e., whenever  $t$  is sufficiently small (cf. (E.2)), the remainder in the Taylor approximation, i.e., the  $O(t^3)$  terms, will be dominated by  $\frac{\mu}{n}(f(\bar{\sigma}) - f(\sigma^k))$ . In particular, if  $\sigma^0$  is sufficiently close to  $\bar{\sigma}$  to satisfy  $O(t^3) \geq -\frac{\mu}{2n}(f(\bar{\sigma}) - f(\sigma^k))$  in the above inequality, we then have

$$\|\text{grad}f(\sigma^k)\|_{\mathbb{F}}^2 \geq \frac{\mu}{2n}(f(\bar{\sigma}) - f(\sigma^k)). \quad (\text{E.8})$$

Combining this inequality with (E.1), we get

$$f(\sigma^{k+1}) - f(\sigma^k) \geq \frac{\mu}{4n^2\|\mathbf{A}\|_1}(f(\bar{\sigma}) - f(\sigma^k)). \quad (\text{E.9})$$

Rearranging terms in the above inequality concludes the proof.

In order to quantify how close  $\sigma^0$  and  $\sigma$  should be so that this convergence rate holds, we need to derive explicit bounds on the higher order terms in (E.4) and (E.6), which we do in the following. The Taylor expansion of  $\sigma^k$  around  $\sigma$  yields

$$\begin{aligned} \sigma_i^k &= \sigma_i \cos(\|u_i\|t) + \frac{u_i}{\|u_i\|} \sin(\|u_i\|t), \\ &= \sigma_i \left[ \sum_{\ell=0}^{\infty} \frac{(-1)^\ell}{(2\ell)!} (\|u_i\|t)^{2\ell} \right] + \frac{u_i}{\|u_i\|} \left[ \sum_{\ell=0}^{\infty} \frac{(-1)^\ell}{(2\ell+1)!} (\|u_i\|t)^{2\ell+1} \right]. \end{aligned}$$

Using this expansion, we can compute  $f(\sigma^k) = \sum_{i,j=1}^n A_{ij} \langle \sigma_i^k, \sigma_j^k \rangle$ . The first three terms in the expansion are already given in (E.5) as follows

$$f(\sigma^k) = f(\sigma) + t^2 \sum_{i=1}^n \left( \langle u_i, v_i \rangle - \|u_i\|^2 \|g_i\| \right) + \beta_f, \quad (\text{E.10})$$

where  $\beta_f$  represents the higher order terms. In order to find an upper bound on  $|\beta_f|$ , we use the Cauchy-Schwarz inequality in the higher order terms in the expansion of  $f(\sigma^k)$ , which yields the following bound

$$|\beta_f| \leq \sum_{i,j=1}^n |A_{ij}| \left( \sum_{\ell=3}^{\infty} \frac{t^\ell}{\ell!} (\|u_i\| + \|u_j\|)^\ell \right).$$

As  $\|\mathbf{u}\|_{\mathbb{F}} = 1$ , we have  $\|u_i\| \leq 1$  for all  $i \in [n]$ , which implies

$$|\beta_f| \leq \sum_{i,j=1}^n |A_{ij}| \left( \sum_{\ell=3}^{\infty} \frac{t^\ell}{\ell!} 2^\ell \right),$$

where we note that  $t$  denotes the geodesic distance between  $\sigma^k$  and  $[\bar{\sigma}]$  as highlighted in (E.2). Assuming that  $t \leq 1$ , we obtain the following upper bound

$$|\beta_f| \leq t^3 n \|\mathbf{A}\|_1 \left( \sum_{\ell=3}^{\infty} \frac{2^\ell}{\ell!} \right).$$

Using the inequality  $\sum_{\ell=3}^{\infty} \frac{2^\ell}{\ell!} = e^2 - 5 \leq 5/2$  above, we get

$$|\beta_f| \leq \frac{5n\|\mathbf{A}\|_1 t^3}{2}.$$

Plugging this value back in (E.10), we obtain

$$f(\boldsymbol{\sigma}^k) \leq f(\boldsymbol{\sigma}) + t^2 \sum_{i=1}^n \left( \langle u_i, v_i \rangle - \|u_i\|^2 \|g_i\| \right) + \frac{5n \|\mathbf{A}\|_1 t^3}{2}. \quad (\text{E.11})$$

Considering the same expansion for  $\|\text{grad} f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 = 2 \sum_{i=1}^n (\|g_i^k\|^2 - \langle \sigma_i^k, g_i^k \rangle^2)$ , we get the following (see (E.4)):

$$\|\text{grad} f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 = 2t^2 \sum_{i=1}^n \left( \|u_i\| \|g_i\| - \left\langle \frac{u_i}{\|u_i\|}, v_i \right\rangle \right)^2 + \beta_g, \quad (\text{E.12})$$

where  $\beta_g$  represents the higher order terms. Upper bounding each higher order terms using the Cauchy-Schwarz inequality as follows, we obtain

$$|\beta_g| \leq 2 \sum_{i=1}^n \left[ \sum_{j,m=1}^n |A_{ij}| |A_{im}| \left( \sum_{\ell=3}^{\infty} \frac{t^\ell}{\ell!} (\|u_j\| + \|u_m\|)^\ell \right) + \sum_{j,m=1}^n |A_{ij}| |A_{im}| \left( \sum_{\substack{\ell,s=0 \\ \ell+s \geq 3}}^{\infty} \frac{t^{\ell+s}}{\ell! s!} (\|u_i\| + \|u_j\|)^{\ell+s} \right) \right].$$

Using the fact that  $\|u_i\| \leq 1$  for all  $i \in [n]$ , we get the following upper bound

$$|\beta_g| \leq 2 \sum_{i=1}^n \left[ \sum_{j,m=1}^n |A_{ij}| |A_{im}| \left( \sum_{\ell=3}^{\infty} \frac{t^\ell}{\ell!} 2^\ell \right) + \sum_{j,m=1}^n |A_{ij}| |A_{im}| \left( \sum_{\substack{\ell,s=0 \\ \ell+s \geq 3}}^{\infty} \frac{t^{\ell+s}}{\ell! s!} 2^{\ell+s} \right) \right].$$

Using the upper bound  $\sum_{j,m=1}^n |A_{ij}| |A_{im}| \leq \|\mathbf{A}\|_1^2$  above, we obtain

$$|\beta_g| \leq 2 \|\mathbf{A}\|_1^2 \sum_{i=1}^n \left[ \sum_{\ell=3}^{\infty} \frac{t^\ell}{\ell!} 2^\ell + \sum_{\substack{\ell,s=0 \\ \ell+s \geq 3}}^{\infty} \frac{t^{\ell+s}}{\ell! s!} 2^{\ell+s} \right].$$

Introducing a change of variables in the last sum, we get

$$\begin{aligned} |\beta_g| &\leq 2 \|\mathbf{A}\|_1^2 \sum_{i=1}^n \left[ \sum_{\ell=3}^{\infty} \frac{t^\ell}{\ell!} 2^\ell + \sum_{\ell=3}^{\infty} \frac{t^\ell}{\ell!} 2^\ell \left( \sum_{s=0}^{\ell} \frac{\ell!}{s! (\ell-s)!} \right) \right], \\ &= 2 \|\mathbf{A}\|_1^2 \sum_{i=1}^n \left[ \sum_{\ell=3}^{\infty} \frac{t^\ell}{\ell!} (2^\ell + 4^\ell) \right]. \end{aligned}$$

Assuming that  $t \leq 1$ , we obtain the following upper bound

$$|\beta_g| \leq 2 \|\mathbf{A}\|_1^2 t^3 \sum_{i=1}^n \left[ \sum_{\ell=3}^{\infty} \frac{1}{\ell!} (2^\ell + 4^\ell) \right].$$

Using the inequality  $\sum_{\ell=3}^{\infty} \frac{2^\ell + 4^\ell}{\ell!} = e^2 + e^4 - 18 \leq 44$  above, we get

$$|\beta_g| \leq 88n \|\mathbf{A}\|_1^2 t^3.$$

Plugging this value back in (E.12), we obtain

$$\|\text{grad} f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 \geq 2t^2 \sum_{i=1}^n \left( \|u_i\| \|g_i\| - \left\langle \frac{u_i}{\|u_i\|}, v_i \right\rangle \right)^2 - 88n \|\mathbf{A}\|_1^2 t^3. \quad (\text{E.13})$$

Using the same bounding technique as in (E.7), we get

$$\begin{aligned}\|\text{grad}f(\boldsymbol{\sigma}^k)\|_F^2 &\geq \frac{\mu}{n} \left( f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k) - \frac{5n\|\mathbf{A}\|_1 t^3}{2} \right) - 88n\|\mathbf{A}\|_1^2 t^3, \\ &= \frac{\mu}{n} \left( f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k) \right) - t^3\|\mathbf{A}\|_1(3\mu + 88n\|\mathbf{A}\|_1).\end{aligned}$$

Therefore, in order for (E.8) to hold, we need

$$t^3\|\mathbf{A}\|_1(3\mu + 88n\|\mathbf{A}\|_1) \leq \frac{\mu}{2n} \left( f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k) \right),$$

which can be equivalently rewritten as follows

$$t^3 \leq \frac{\mu(f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k))}{2n\|\mathbf{A}\|_1(3\mu + 88n\|\mathbf{A}\|_1)}.$$

As  $f(\boldsymbol{\sigma}^k)$  is a monotonically non-decreasing sequence, then as soon as  $\boldsymbol{\sigma}^0$  is sufficiently close to  $[\bar{\boldsymbol{\sigma}}]$  in the sense that

$$\text{dist}(\boldsymbol{\sigma}^0, [\bar{\boldsymbol{\sigma}}]) \leq \left( \frac{\mu(f(\bar{\boldsymbol{\sigma}}) - f(\boldsymbol{\sigma}^k))}{2n\|\mathbf{A}\|_1(3\mu + 88n\|\mathbf{A}\|_1)} \right)^{1/3},$$

then the linear convergence rate presented in (E.9) holds.

## F Proof of Theorem 6

Suppose  $\mathbf{X}^* = \boldsymbol{\sigma}\boldsymbol{\sigma}^\top$ , where  $\boldsymbol{\sigma} \in \mathcal{M}_r$  and  $\text{rank}(\mathbf{X}^*) = r^* \leq r$ . We have  $\mathbf{Z}^* = \boldsymbol{\Lambda} - \mathbf{A}$  by the definition of the dual problem, where  $\boldsymbol{\Lambda} = \text{Diag}(\|g_1\|, \dots, \|g_n\|)$ . Due to strict complementarity, we have  $\text{rank}(\mathbf{Z}^*) = n - r^*$  and kernel of  $\mathbf{Z}^*$  is equal to the column space of  $\mathbf{X}^*$ , i.e.,  $\ker(\mathbf{Z}^*) = \text{col}(\mathbf{X}^*)$ . Since  $\mathbf{X}^* = \boldsymbol{\sigma}\boldsymbol{\sigma}^\top$  and  $\mathbf{Z}^* = \boldsymbol{\Lambda} - \mathbf{A}$ , we equivalently have  $\ker(\boldsymbol{\Lambda} - \mathbf{A}) = \text{col}(\boldsymbol{\sigma})$ . As  $\mathbf{Z}^*$  is feasible for the dual, then  $\mathbf{Z}^* = \boldsymbol{\Lambda} - \mathbf{A} \succeq 0$ , and consequently  $\langle \mathbf{u}, (\boldsymbol{\Lambda} - \mathbf{A})\mathbf{u} \rangle \geq 0$ , for all  $\mathbf{u} \in \mathbb{R}^{n \times r}$ .

Now consider the quadratic form  $h(\mathbf{u}) := \langle \mathbf{u}, (\boldsymbol{\Lambda} - \mathbf{A})\mathbf{u} \rangle$  over  $\mathbf{u} \in T_{\boldsymbol{\sigma}}\mathcal{M}_r$ . First, we show that  $h(\mathbf{u}) = 0$  if and only if  $\mathbf{u} \in \mathcal{V}_{\boldsymbol{\sigma}}$ . The *only if* direction of the proof is straightforward, i.e.,  $(\boldsymbol{\Lambda} - \mathbf{A})\boldsymbol{\sigma} = 0$  and  $\mathbf{u} = \boldsymbol{\sigma}\mathbf{B}$  for some skew-symmetric matrix  $\mathbf{B}$  directly imply  $h(\mathbf{u}) = 0$  for all  $\mathbf{u} \in \mathcal{V}_{\boldsymbol{\sigma}}$ . To show the *if* direction, let  $\mathbf{u} \in T_{\boldsymbol{\sigma}}\mathcal{M}_r$  such that  $h(\mathbf{u}) = 0$ , or equivalently  $\text{tr}((\boldsymbol{\Lambda} - \mathbf{A})\mathbf{u}\mathbf{u}^\top) = 0$ . As both  $\boldsymbol{\Lambda} - \mathbf{A}$  and  $\mathbf{u}\mathbf{u}^\top$  are positive semidefinite matrices, this implies  $(\boldsymbol{\Lambda} - \mathbf{A})\mathbf{u} = 0$ . Therefore, columns of  $\mathbf{u}$  are in  $\ker(\boldsymbol{\Lambda} - \mathbf{A}) = \text{col}(\boldsymbol{\sigma})$ , which implies there exists  $\mathbf{B} \in \mathbb{R}^{r \times r}$  such that  $\mathbf{u} = \boldsymbol{\sigma}\mathbf{B}$ . As  $\mathbf{u} \in T_{\boldsymbol{\sigma}}\mathcal{M}_r$ , then  $\langle \sigma_i, u_i \rangle = \langle \sigma_i, \mathbf{B}^\top \sigma_i \rangle = \langle \sigma_i \sigma_i^\top, \mathbf{B} \rangle = 0$ , for all  $i \in [n]$ . Without loss of generality, assume that the last  $r - r^*$  columns of  $\boldsymbol{\sigma}$  are equal to zero. Then, by dual nondegeneracy of the SDP, the principal submatrices of dimension  $r^* \times r^*$  of  $\{\sigma_i \sigma_i^\top\}_{i=1}^n$  spans  $\mathcal{S}^{r^*}$ . Consider the decomposition

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{bmatrix},$$

where  $\mathbf{B}_{11} \in \mathbb{R}^{r^* \times r^*}$  and  $\mathbf{B}_{22} \in \mathbb{R}^{(r-r^*) \times (r-r^*)}$ . Then, the dual nondegeneracy implies that  $\mathbf{B}_{11}$  is a skew-symmetric matrix, i.e.,  $\mathbf{B}_{11}^\top = -\mathbf{B}_{11}$ . Furthermore, as the last  $r - r^*$  columns of  $\boldsymbol{\sigma}$  are equal to zero, then  $\mathbf{u} = \boldsymbol{\sigma}\mathbf{B}$  does not depend on  $\mathbf{B}_{21}$  and  $\mathbf{B}_{22}$ . Therefore, we can pick  $\mathbf{B}_{21} = -\mathbf{B}_{12}^\top$  and  $\mathbf{B}_{22} = 0$  such that  $\mathbf{B}$  is a skew-symmetric matrix and observe that  $\mathbf{u} \in \mathcal{V}_{\boldsymbol{\sigma}}$ .

To conclude the proof, we let  $\{\mathbf{u}^\ell\}_{\ell=1}^{n(r-1)}$  be an orthogonal basis to  $T_{\boldsymbol{\sigma}}\mathcal{M}_r$  such that  $\{\mathbf{u}^\ell\}_{\ell=1}^s$  is a basis for  $\mathcal{V}_{\boldsymbol{\sigma}}$ . Let  $\bar{\mathbf{Q}} \in \mathbb{R}^{n(r-1) \times n(r-1)}$  such that  $\bar{Q}_{ij} = \langle \mathbf{u}^i, (\boldsymbol{\Lambda} - \mathbf{A})\mathbf{u}^j \rangle$ . Consider the function

$\bar{h} : \mathbb{R}^{n(r-1)} \rightarrow \mathbb{R}^{n(r-1)}$  such that  $\bar{h}(\mathbf{v}) = \mathbf{v}^\top \bar{\mathbf{Q}} \mathbf{v}$  and observe that  $\bar{h}(\text{vec}(\mathbf{u})) = h(\mathbf{u})$ . Let  $\mathbf{L} = [\text{vec}(\mathbf{u}^1), \dots, \text{vec}(\mathbf{u}^s)] \in \mathbb{R}^{n(r-1) \times s}$ , then  $\mathbf{v}^\top \bar{\mathbf{Q}} \mathbf{v} < 0$  for all  $\mathbf{v} \in \ker(\mathbf{A}) \setminus \{\mathbf{0}\}$ . Then, by Finsler's Lemma,  $(\mathbf{L}^\perp)^\top \bar{\mathbf{Q}} \mathbf{L}^\perp \succ 0$ . Equivalently, there exists  $\mu > 0$  such that  $h(\mathbf{u}) \geq \mu \|\mathbf{u}\|_{\mathbb{F}}^2$  for all  $\mathbf{u} \in T_{\sigma} \mathcal{M}_r \setminus \mathcal{V}_{\sigma}$ .

## G An Ascent Lemma for the Second-Order Oracle

**Lemma 14.** *Let  $\mathbf{u}^k \in T_{\sigma^k} \mathcal{M}_r$  be a tangent vector of  $\sigma^k$  such that  $\|\mathbf{u}^k\|_{\mathbb{F}} = 1$ ,  $\langle \mathbf{u}^k, \text{grad} f(\sigma^k) \rangle \geq 0$ , and  $\langle \mathbf{u}^k, \text{Hess} f(\sigma^k)[\mathbf{u}^k] \rangle \geq \varepsilon/2$ . Consider the update rule given by the exponential map  $\sigma^{k+1} = \text{Exp}_{\sigma^k}(\mathbf{u}^k)$ , i.e.,*

$$\sigma_i^{k+1} = \sigma_i^k \cos(\|u_i^k\|t) + \frac{u_i^k}{\|u_i^k\|} \sin(\|u_i^k\|t), \quad \text{for all } i \in [n],$$

where  $t = \frac{\varepsilon}{15\|\mathbf{A}\|_1}$  is the step size. These iterates satisfy the following ascent in the function value:

$$f(\sigma^{k+1}) - f(\sigma^k) \geq \frac{\varepsilon^3}{2700\|\mathbf{A}\|_1^2}.$$

**Proof** The second-order oracle returns a vector  $\mathbf{u}^k \in T_{\sigma^k} \mathcal{M}_r$  such that  $\|\mathbf{u}^k\|_{\mathbb{F}} = 1$ ,  $\langle \mathbf{u}^k, \text{grad} f(\sigma^k) \rangle \geq 0$ , and  $\langle \mathbf{u}^k, \text{Hess} f(\sigma^k)[\mathbf{u}^k] \rangle \geq \varepsilon/2$ . Then we set the next iterate  $\sigma^{k+1}$  as follows

$$\sigma_i^{k+1} = \sigma_i^k \cos(\|u_i^k\|t) + \frac{u_i^k}{\|u_i^k\|} \sin(\|u_i^k\|t), \quad (\text{G.1})$$

where  $0 < t < 1$  can be viewed as step size. The third order Taylor approximation to (G.1) yields:

$$\sigma_i^{k+1} = \sigma_i^k + tu_i^k - \frac{t^2}{2} \|u_i^k\|^2 \sigma_i - \frac{t^3}{6} \|u_i^k\|^2 u_i^k + O(t^4),$$

and using this approximation, we obtain

$$g_i^{k+1} = g_i^k + tv_i^k - \frac{t^2}{2} \tilde{g}_i^k - \frac{t^3}{6} \tilde{v}_i^k + O(t^4),$$

where

$$v_i^k = \sum_{j \neq i} A_{ij} u_j^k \quad \text{and} \quad \tilde{g}_i^k = \sum_{j \neq i} A_{ij} \|u_j^k\|^2 \sigma_j^k \quad \text{and} \quad \tilde{v}_i^k = \sum_{j \neq i} A_{ij} \|u_j^k\|^2 u_j^k. \quad (\text{G.2})$$

This yields the following Taylor approximation to  $f(\sigma^{k+1})$ , which can be written as follows

$$\begin{aligned} f(\sigma^{k+1}) &= \sum_{i=1}^n \langle \sigma_i^{k+1}, g_i^{k+1} \rangle, \\ &= \sum_{i=1}^n \langle \sigma_i^k + tu_i^k - \frac{t^2}{2} \|u_i^k\|^2 \sigma_i - \frac{t^3}{6} \|u_i^k\|^2 u_i^k, g_i^k + tv_i^k - \frac{t^2}{2} \tilde{g}_i^k - \frac{t^3}{6} \tilde{v}_i^k \rangle + O(t^4), \\ &= \sum_{i=1}^n \langle \sigma_i^k, g_i^k \rangle + t \sum_{i=1}^n (\langle \sigma_i^k, v_i^k \rangle + \langle u_i^k, g_i^k \rangle) + \frac{t^2}{2} \sum_{i=1}^n (2\langle u_i^k, v_i^k \rangle - \langle \sigma_i^k, \tilde{g}_i^k \rangle - \|u_i^k\|^2 \langle \sigma_i^k, g_i^k \rangle), \\ &\quad - \frac{t^3}{6} \sum_{i=1}^n (\langle \sigma_i^k, \tilde{v}_i^k \rangle + 3\langle u_i^k, \tilde{g}_i^k \rangle + 3\|u_i^k\|^2 \langle \sigma_i^k, v_i^k \rangle + \|u_i^k\|^2 \langle u_i^k, g_i^k \rangle) + O(t^4). \end{aligned}$$



Plugging in the definitions in (G.2), we obtain

$$\begin{aligned} f(\boldsymbol{\sigma}^{k+1}) &= \sum_{i=1}^n \langle \boldsymbol{\sigma}_i^k, \mathbf{g}_i^k \rangle + t \sum_{i=1}^n 2 \langle \mathbf{u}_i^k, \mathbf{g}_i^k \rangle + \frac{t^2}{2} \sum_{i=1}^n 2 \left( \langle \mathbf{u}_i^k, \mathbf{v}_i^k \rangle - \|\mathbf{u}_i^k\|^2 \langle \boldsymbol{\sigma}_i^k, \mathbf{g}_i^k \rangle \right), \\ &\quad - \frac{t^3}{6} \sum_{i=1}^n 4 \|\mathbf{u}_i^k\|^2 \left( \langle \boldsymbol{\sigma}_i^k, \mathbf{v}_i^k \rangle + \langle \mathbf{u}_i^k, \mathbf{g}_i^k \rangle \right) + O(t^4). \end{aligned}$$

Using the definitions of  $f(\boldsymbol{\sigma}^k)$  and its derivatives, the above equality can be written as follows

$$\begin{aligned} f(\boldsymbol{\sigma}^{k+1}) &= f(\boldsymbol{\sigma}^k) + t \langle \mathbf{u}^k, \text{grad} f(\boldsymbol{\sigma}^k) \rangle + \frac{t^2}{2} \langle \mathbf{u}^k, \text{Hess} f(\boldsymbol{\sigma}^k)[\mathbf{u}^k] \rangle \\ &\quad - \frac{t^3}{6} \sum_{i=1}^n 4 \|\mathbf{u}_i^k\|^2 \left( \langle \boldsymbol{\sigma}_i^k, \mathbf{v}_i^k \rangle + \langle \mathbf{u}_i^k, \mathbf{g}_i^k \rangle \right) + O(t^4). \end{aligned}$$

Here, our aim is to lower bound the remainder term corresponding to the third and higher order terms. To this end, we upper bound the third order terms using the inequalities  $\langle \boldsymbol{\sigma}_i^k, \mathbf{v}_i^k \rangle \leq \|\mathbf{A}\|_1$  and  $\langle \mathbf{u}_i^k, \mathbf{g}_i^k \rangle \leq \|\mathbf{A}\|_1$ . This yields

$$\sum_{i=1}^n 4 \|\mathbf{u}_i^k\|^2 \left( \langle \boldsymbol{\sigma}_i^k, \mathbf{v}_i^k \rangle + \langle \mathbf{u}_i^k, \mathbf{g}_i^k \rangle \right) \leq 8 \|\mathbf{A}\|_1 \sum_{i=1}^n \|\mathbf{u}_i^k\|^2 = 8 \|\mathbf{A}\|_1,$$

where the last equality follows since  $\|\mathbf{u}^k\|_{\text{F}} = 1$ . Considering the structure of the higher order terms and using similar bounds, we observe that each individual inner product term in higher order expansions can be upper bounded by  $\|\mathbf{u}_i\|^{2\lfloor t/2 \rfloor} \|\mathbf{A}\|_1$  and consequently we get the following lower bound:

$$f(\boldsymbol{\sigma}^{k+1}) \geq f(\boldsymbol{\sigma}^k) + t \langle \mathbf{u}^k, \text{grad} f(\boldsymbol{\sigma}^k) \rangle + \frac{t^2}{2} \langle \mathbf{u}^k, \text{Hess} f(\boldsymbol{\sigma}^k)[\mathbf{u}^k] \rangle - t^3 \|\mathbf{A}\|_1 \sum_{\ell=3}^{\infty} \frac{2^\ell}{\ell!}.$$

Using  $\sum_{\ell=3}^{\infty} \frac{2^\ell}{\ell!} = e^2 - 5 \leq 5/2$ , we obtain

$$f(\boldsymbol{\sigma}^{k+1}) \geq f(\boldsymbol{\sigma}^k) + t \langle \mathbf{u}^k, \text{grad} f(\boldsymbol{\sigma}^k) \rangle + \frac{t^2}{2} \langle \mathbf{u}^k, \text{Hess} f(\boldsymbol{\sigma}^k)[\mathbf{u}^k] \rangle - \frac{5 \|\mathbf{A}\|_1}{2} t^3. \quad (\text{G.3})$$

Since we are given that  $\langle \mathbf{u}^k, \text{grad} f(\boldsymbol{\sigma}^k) \rangle \geq 0$  and  $\langle \mathbf{u}^k, \text{Hess} f(\boldsymbol{\sigma}^k)[\mathbf{u}^k] \rangle \geq \varepsilon/2$ , (G.3) yields

$$f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) \geq \frac{\varepsilon}{4} t^2 - \frac{5 \|\mathbf{A}\|_1}{2} t^3.$$

Choosing  $t = \frac{\varepsilon}{15 \|\mathbf{A}\|_1}$  maximizes the right-hand side of the above inequality and guarantees the following ascent in the function value:

$$f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) \geq \frac{\varepsilon^3}{2700 \|\mathbf{A}\|_1^2}.$$

□

## H Proof of Theorem 10

We have already proved that each iteration of BCM yields following improvement (by (C.3)):

$$f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) \geq \frac{\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2}{2n\|\mathbf{A}\|_1} \geq \frac{\varepsilon^3}{2700n\|\mathbf{A}\|_1^2}, \quad (\text{H.1})$$

where the last inequality holds since BCM is applied at iteration  $k$  only if  $\|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 \geq \frac{\varepsilon^3}{1350\|\mathbf{A}\|_1}$ . Similarly, each iteration of the second-order oracle yields the following improvement (by Lemma 14):

$$f(\boldsymbol{\sigma}^{k+1}) - f(\boldsymbol{\sigma}^k) \geq \frac{\varepsilon^3}{2700\|\mathbf{A}\|_1^2}. \quad (\text{H.2})$$

Hence, an epoch ( $n$  iterations) of BCM yields the same amount of function value improvement as an iteration of the second-order oracle. Let

$$f^* = \left(1 - \frac{1}{r-1}\right) \text{SDP}(\mathbf{A})$$

denote the desired approximation ratio and consider the approximation gap of the solution  $\boldsymbol{\sigma}$  with respect to  $f^*$  that is given by

$$h(\boldsymbol{\sigma}) = f^* - f(\boldsymbol{\sigma}). \quad (\text{H.3})$$

The aim of the algorithm is to find a solution  $\boldsymbol{\sigma}$  that satisfy  $h(\boldsymbol{\sigma}) \leq \varepsilon$  for some  $\varepsilon > 0$ . Consider that the BCM2 algorithm runs  $K_{\text{BCM}}$  epochs of BCM and  $K_{\text{H}}$  iterations of the second-order oracle such that a total of  $K = nK_{\text{BCM}} + K_{\text{H}}$  iterations are made. Let  $\mathcal{G} = \{0 \leq k \leq K-1 : \|\text{grad}f(\boldsymbol{\sigma}^k)\|_{\mathbb{F}}^2 \geq \frac{\varepsilon^3}{1350\|\mathbf{A}\|_1}\}$  be the set of iterations at which BCM step is taken and let  $\mathcal{H} = \{0 \leq k \leq K-1\} \setminus \mathcal{G}$  be the set of iterations at which a second-order oracle step is taken. Then, the approximation gap decreases at each iteration by the following amount:

$$h(\boldsymbol{\sigma}^k) - h(\boldsymbol{\sigma}^{k+1}) \geq \frac{\varepsilon^3}{2700\|\mathbf{A}\|_1^2} \delta_k, \quad (\text{H.4})$$

where, for notational simplicity, we introduced

$$\delta_k = \begin{cases} \frac{1}{n}, & \text{if } k \in \mathcal{G}, \\ 1, & \text{if } k \in \mathcal{H}. \end{cases} \quad (\text{H.5})$$

By Theorem 8, we are given that any  $\varepsilon$ -approximate concave point  $\boldsymbol{\sigma}$  satisfies

$$h(\boldsymbol{\sigma}) \leq \frac{n}{2}\varepsilon. \quad (\text{H.6})$$

Hence, the right-hand side of (H.4) can be lower bounded as follows

$$h(\boldsymbol{\sigma}^k) - h(\boldsymbol{\sigma}^{k+1}) \geq \frac{2\delta_k}{675n^3\|\mathbf{A}\|_1^2} h^3(\boldsymbol{\sigma}^k). \quad (\text{H.7})$$

Considering the reciprocal of the approximation gap, we observe that

$$\begin{aligned} \frac{1}{h^2(\boldsymbol{\sigma}^{k+1})} - \frac{1}{h^2(\boldsymbol{\sigma}^k)} &= \frac{(h(\boldsymbol{\sigma}^k) - h(\boldsymbol{\sigma}^{k+1})) (h(\boldsymbol{\sigma}^k) + h(\boldsymbol{\sigma}^{k+1}))}{h^2(\boldsymbol{\sigma}^{k+1})h^2(\boldsymbol{\sigma}^k)}, \\ &\geq \frac{2\delta_k}{675n^3\|\mathbf{A}\|_1^2} \frac{h(\boldsymbol{\sigma}^k) (h(\boldsymbol{\sigma}^k) + h(\boldsymbol{\sigma}^{k+1}))}{h^2(\boldsymbol{\sigma}^{k+1})}, \end{aligned} \quad (\text{H.8})$$

where the inequality follows by (H.7). As the right-hand side of (H.7) is lower bounded by zero, we have  $h(\boldsymbol{\sigma}^k) \geq h(\boldsymbol{\sigma}^{k+1})$ . Thus, we can lower bound the right-hand side of (H.8) as follows

$$\frac{1}{h^2(\boldsymbol{\sigma}^{k+1})} - \frac{1}{h^2(\boldsymbol{\sigma}^k)} \geq \frac{4\delta_k}{675n^3\|\mathbf{A}\|_1^2}. \quad (\text{H.9})$$

Summing (H.9) over  $k = 0, 1, \dots, K-1$ , we get

$$\frac{1}{h^2(\boldsymbol{\sigma}^K)} - \frac{1}{h^2(\boldsymbol{\sigma}^0)} \geq \sum_{k=0}^{K-1} \frac{4\delta_k}{675n^3\|\mathbf{A}\|_1^2} = \frac{4}{675n^3\|\mathbf{A}\|_1^2}(K_{\text{BCM}} + K_{\text{H}}).$$

Given that  $\boldsymbol{\sigma}^0$  is not an  $\varepsilon$ -approximate concave point (or else, there is nothing to prove), we have

$$\frac{1}{h^2(\boldsymbol{\sigma}^K)} \geq \frac{4}{675n^3\|\mathbf{A}\|_1^2}(K_{\text{BCM}} + K_{\text{H}}). \quad (\text{H.10})$$

Since by (H.6), we know that  $\frac{1}{h(\boldsymbol{\sigma})} \geq \frac{2}{n\varepsilon}$  for any  $\varepsilon$ -approximate concave point, then as soon as

$$K_{\text{BCM}} + K_{\text{H}} \geq \frac{675n\|\mathbf{A}\|_1^2}{\varepsilon^2} \quad (\text{H.11})$$

iterations made, BCM2 is guaranteed to return an  $\varepsilon$ -approximate concave point, i.e., there exists a solution  $\boldsymbol{\sigma}^k$  for some  $1 < k < K$  such that  $h(\boldsymbol{\sigma}^k) \leq \frac{n}{2}\varepsilon$ . Since  $\{h(\boldsymbol{\sigma}^k)\}_{k \geq 0}$  is a nonincreasing sequence (as we have already shown in (H.7)), then the final iterate of the algorithm  $\boldsymbol{\sigma}^K$  is guaranteed to satisfy  $h(\boldsymbol{\sigma}^K) \leq \frac{n}{2}\varepsilon$ , i.e.,  $\boldsymbol{\sigma}^K$  is an  $\varepsilon$ -approximate concave point.

## I Proof of Theorem 12

**Theorem 15** ([KW92, Theorem 4.2]). *Let  $\mathbf{A} \in \mathbb{R}^{n \times n}$  be a positive semidefinite matrix,  $b \in \mathbb{R}^n$  be an arbitrary vector and  $\lambda_L^\ell(\mathbf{A}, b)$  denote the output of the Lanczos algorithm after  $\ell$  iterations when applied to find the leading eigenvalue of  $\mathbf{A}$  (denoted by  $\lambda_1(\mathbf{A})$ ) with initialization  $b$ . In particular,*

$$\lambda_L^\ell(\mathbf{A}, b) = \max \left\{ \frac{\langle x, \mathbf{A}x \rangle}{\langle x, x \rangle} : 0 \neq x \in \text{span}(b, \dots, \mathbf{A}^{\ell-1}b) \right\}.$$

*Assume that  $b$  is uniformly distributed over the set  $\{b \in \mathbb{R}^n : \|b\| = 1\}$  and let  $\epsilon \in [0, 1)$ . Then, the probability that the Lanczos algorithm does not return an  $\epsilon$ -approximation to the leading eigenvalue of  $\mathbf{A}$  exponentially decreases as follows*

$$\mathbb{P}\left(\lambda_L^\ell(\mathbf{A}, b) < (1 - \epsilon)\lambda_1(\mathbf{A})\right) \begin{cases} \leq 1.648\sqrt{n}e^{-\sqrt{\epsilon}(2\ell-1)}, & \text{if } 0 < \ell < n(r-1), \\ = 0, & \text{if } \ell \geq n(r-1). \end{cases}$$

Since the tangent space  $T_{\boldsymbol{\sigma}}\mathcal{M}_r$  has dimension  $n(r-1)$ , then we can define a symmetric matrix (where we drop the notational dependency on  $\boldsymbol{\sigma}$  for simplicity)  $\mathbf{H} \in \mathbb{R}^{n(r-1) \times n(r-1)}$  that represents the linear operator  $\text{Hess}f(\boldsymbol{\sigma})$  in the basis  $\{\mathbf{u}^1, \dots, \mathbf{u}^{n(r-1)}\}$  such that  $\text{span}(\mathbf{u}^1, \dots, \mathbf{u}^{n(r-1)}) = T_{\boldsymbol{\sigma}}\mathcal{M}_r$ . In particular, letting  $H_{ij} = \langle \mathbf{u}^i, \text{Hess}f(\boldsymbol{\sigma})[\mathbf{u}^j] \rangle$  yields the desired matrix  $\mathbf{H}$  and the Lanczos algorithm is run to find the leading eigenvalue of this matrix. Here, it is important to note that  $\mathbf{H}$  is not a psd matrix, so it is required to shift  $\mathbf{H}$  with a large enough multiple of the identity matrix so that the resulting matrix is guaranteed to be positive semidefinite. In particular, by inspecting the definition of  $\text{Hess}f(\boldsymbol{\sigma})$  in (A.2), it is easy to observe that  $\|\text{Hess}f(\boldsymbol{\sigma})\|_{\text{op}} \leq 4\|\mathbf{A}\|_1$ . Therefore,

it is sufficient to run the Lanczos algorithm to find the leading eigenvalue of  $\widetilde{\mathbf{H}} = \mathbf{H} + 4\|\mathbf{A}\|_1\mathbf{I}$ , where  $\mathbf{I}$  denotes the appropriate sized identity matrix. On the other hand, we initialize the Lanczos algorithm with a random vector  $\mathbf{u}$  of unit norm (i.e.,  $\|\mathbf{u}\|_F = 1$ ) in the tangent space  $T_{\sigma}\mathcal{M}_r$ . Notice that  $\mathbf{u}$  can equivalently be represented as a vector  $b \in \mathbb{R}^{n(r-1)}$  in the basis  $\{\mathbf{u}^1, \dots, \mathbf{u}^{n(r-1)}\}$  as  $\mathbf{u} = \sum_{i=1}^{n(r-1)} b_i \mathbf{u}^i$  such that  $\|b\| = 1$ . Then, by Theorem 15, we have

$$\mathbb{P}\left(\lambda_L^\ell(\widetilde{\mathbf{H}}, b) < (1 - \epsilon)\lambda_1(\widetilde{\mathbf{H}})\right) \leq 1.648\sqrt{n(r-1)}e^{-\sqrt{\epsilon}(2\ell-1)}.$$

Letting  $\lambda_1(\mathbf{H})$  denote the leading eigenvalue of  $\mathbf{H}$ , we run the Lanczos algorithm to obtain a vector  $b^*$  such that  $\|b^*\| = 1$  and  $\langle b^*, \mathbf{H}b^* \rangle \geq \lambda_1(\mathbf{H})/2$ . Thus, we want  $\mathbb{P}\left(\lambda_L^\ell(\widetilde{\mathbf{H}}, b) < 4\|\mathbf{A}\|_1 + \lambda_1(\mathbf{H})/2\right)$  to be small. Setting  $\epsilon^* = \frac{\lambda_1(\mathbf{H})}{16\|\mathbf{A}\|_1}$ , we can observe that

$$\begin{aligned} (1 - \epsilon^*)\lambda_1(\widetilde{\mathbf{H}}) &= \left(1 - \frac{\lambda_1(\mathbf{H})}{16\|\mathbf{A}\|_1}\right)(4\|\mathbf{A}\|_1 + \lambda_1(\mathbf{H})), \\ &= 4\|\mathbf{A}\|_1 + \frac{3\lambda_1(\mathbf{H})}{4} - \frac{(\lambda_1(\mathbf{H}))^2}{16\|\mathbf{A}\|_1}, \\ &\geq 4\|\mathbf{A}\|_1 + \frac{\lambda_1(\mathbf{H})}{2}, \end{aligned}$$

where the inequality follows since  $\lambda_1(\mathbf{H}) \leq 4\|\mathbf{A}\|_1$ . Consequently, we have

$$\mathbb{P}\left(\lambda_L^\ell(\widetilde{\mathbf{H}}, b) < 4\|\mathbf{A}\|_1 + \lambda_1(\mathbf{H})/2\right) \leq \mathbb{P}\left(\lambda_L^\ell(\widetilde{\mathbf{H}}, b) < (1 - \epsilon^*)\lambda_1(\widetilde{\mathbf{H}})\right) \leq 1.648\sqrt{n(r-1)}e^{-\sqrt{\epsilon^*}(2\ell-1)}.$$

By Theorem 10, we know that the Lanczos method is called at most  $\lceil 675n\|\mathbf{A}\|_1^2/\epsilon^2 \rceil$  times to search for an  $\epsilon$ -approximate concave point and for any non-desired solution we have  $\lambda_1(\mathbf{H}) \geq \epsilon$  by the definition of  $\epsilon$ -approximate concave point. Then, by using a union bound over all calls to the Lanczos method, we conclude that when the Lanczos method is run for  $\ell$  iterations, we have the following guarantee

$$\begin{aligned} \mathbb{P}(\text{Algorithm 2+3 fails to return an } \epsilon\text{-approximate concave point}) \\ \leq \left\lceil \frac{675n\|\mathbf{A}\|_1^2}{\epsilon^2} \right\rceil 1.648\sqrt{n(r-1)}e^{-\sqrt{\frac{\epsilon}{16\|\mathbf{A}\|_1}}(2\ell-1)}. \end{aligned}$$

In order to set this probability to some  $\delta \in (0, 1)$ , we let

$$\ell^* = \left\lceil \left( \frac{1}{2} + 2\sqrt{\frac{\|\mathbf{A}\|_1}{\epsilon}} \right) \log \left( \frac{\left\lceil \frac{675n\|\mathbf{A}\|_1^2}{\epsilon^2} \right\rceil 1.648\sqrt{n(r-1)}}{\delta} \right) \right\rceil = \tilde{\mathcal{O}} \left( \sqrt{\frac{\|\mathbf{A}\|_1}{\epsilon}} \log \left( \frac{n\sqrt{n(r-1)}}{\delta} \right) \right),$$

where tilde is used to hide poly-logarithmic factors in  $\|\mathbf{A}\|_1/\epsilon$ . Since the Lanczos algorithm is guaranteed to return the leading eigenvalue with probability 1 in at most  $n(r-1)$  iterations, then running each Lanczos subroutine for  $\min(\ell^*, n(r-1))$  iterations, it is guaranteed that Algorithm 2+3 returns an  $\epsilon$ -approximate concave point with probability at least  $1 - \delta$ .