

The Deflation-Inflation Method for Certain Semidefinite Programming and Maximum Determinant Completion Problems

Leonid Gurvits and Peter N. Yianilos
(Extended Abstract)*

September 21, 1999

Abstract

The *deflation-inflation* convex optimization method is introduced. One result is a simple and practical approximation algorithm for the MAX CUT problem based on the semidefinite programming relaxation and analysis of Goemans and Williamson. Another consequence is a closed-form expression for the maximum-determinant completion of a positive definite band matrix. Local and global convergence results are established, along with other properties. The overall development reveals an interesting interplay between the language and outlook of probability, and that of positive definite matrices and mathematical programming.

Keywords: Semidefinite programming, Maximum cut (MAX CUT) problem, maximum entropy, maximum determinant completion, Bregman distance, Kullback-Leibler distance, normal (Gaussian) density, Sinkhorn iteration.

1 Introduction

Semidefinite programming has emerged as an important and practical optimization framework that is increasingly relevant to combinatorial problems. Given symmetric matrices C and Q_1, \dots, Q_k , and vector b , the objective is to minimize $C \bullet P$ subject to $Q_i \bullet P = b_i, i = 1, \dots, k$ and $P \succeq 0$.

There has been a great deal of recent work on the solution of semidefinite programming problems. The ellipsoid method together with a separation oracle due to Lovász may be used; and the work of Nesterov and Nemirovsky [21] and Alizadeh [2] adapts the interior-point method to the task.

The connection between semidefinite programming and combinatorial optimization was first established by Lovász in [17] and later by Lovász and Schrijver in [18].

Goemans and Williamson [12] discovered a

semidefinite programming relaxation of the NP-hard MAX CUT problem, which generates cuts with weight no less than ≈ 0.87856 of optimal. Other optimization problems have been considered using essentially the same approach.

For MAX CUT their method begins by constructing the symmetric weighted adjacency matrix C for the graph G under consideration. We denote by n the number of vertices in G , and by m the number of edges. The semidefinite programming problem is then to find matrix P that minimizes $C \bullet P$ subject to $P \succ 0$ and $\text{diag}(P) = I$. This matrix is then factored as $BB^t = P$ using, for example, Cholesky decomposition. Given a vector r they interpret the sign pattern of Br as a cut — and show that for uniform random r the expected value of this cut is at least ≈ 0.87856 times the weight of a maximum cut. It is easy to see that instead of uniform random r one can use a Gaussian vector ξ with unit covariance matrix. Also, since only the distribution of the sign pattern matters it is enough to get P to within a diagonal factor, i.e. DPD , where D is a diagonal positive-definite matrix.

This approach and the general connection of semidefinite programming to combinatorial optimization is surveyed in [11]. An alternative solution approach is given by Klein and Lu [15] based on the power-method. Benson Ye and Zhang describe a *dual scaling* algorithm and report on experiments for semidefinite programs arising from instances of MAX CUT [3]. The MAX CUT problem is surveyed in [22].

This paper introduces the *deflation-inflation* algorithm and outlook for certain convex optimization problems including those arising from the MAX CUT relaxation above. Applied to MAX CUT the algorithm is quite simple. It begins by choosing diagonal values for C so that $C \succ 0$. We solve an approximation to the semidefinite program and next choose $\epsilon > 0$ to select the desired precision.

In each cycle of the algorithm the diagonal entries of C are modified one-by-one, i.e. for each iteration only one diagonal entry is changed. The update rule is:

*Both authors are with the NEC Research Institute, 4 Independence Way, Princeton NJ, 08540. Email: (gurvits/pny)@research.nj.nec.com

$$C_{i,i} := C_{i,i} - (C_{i,i}^{-1})^{-1} + \epsilon$$

The algorithm terminates when each $C_{i,i}^{-1}$ is equal to ϵ within some error tolerance. Notice that the algorithm is essentially searching over the space of diagonal values for C .

We show that the algorithm above converges to a unique solution, and that this approximates well the original problem. Our analysis further establishes local linear convergence under rather general circumstances – and global convergence after $O(1/\epsilon \log 1/\epsilon)$ diagonal updates given a form of the algorithm that chooses diagonal elements *greedily*. The local analysis and portions of the global analysis apply to the well-studied Sinkhorn balancing problem [23] and seems relevant to the recent matching work of Linial, Samorodnitsky and Widgerson [20]. Our global analysis is the first we are aware of for projection methods of this general class, which includes the celebrated *expectation maximization* (EM) algorithm [9].

The critical operation in each update is $(C_{i,i}^{-1})^{-1}$. Notice that only one element of C^{-1} is examined. The operation then amounts to solving a system of linear equations. So when C has special structure, or is sparse, specialized statements about the algorithm’s complexity can be made. When these systems are large and can be rapidly solved, our method may be the best available.

In the sparse case it is interesting to note that our basic algorithm is linear space, i.e. $O(m)$.

We have implemented our algorithm for MAX CUT using Gauss-Seidel iteration with successive over-relaxation to solve the linear systems [8]. To avoid inverting the resulting P and performing Cholesky decomposition, we instead compute $P^{-1/2}r$ as a power series. For sparse constant degree random graphs our analysis suggests that this overall method should scale roughly as $O(n^2)$ and our experiments confirm this. The full paper describes this implementation and our experimental results.

Recall that the starting matrix C is essentially a combinatorial object representing the weighted adjacency of graph nodes. After the algorithm runs, modifying only its diagonal entries, C^{-1} is the matrix P we seek. An interesting aspect of this work is that a probabilistic viewpoint is of utility when addressing this essentially combinatorial problem. We view P as the covariance matrix associated with a normal density, and C gives its inverse covariance. It is easily verified that the random cuts are then just sign patterns from random vectors drawn from this density.

Each step in our algorithm corresponds to a projection operation with respect to Kullback-Leibler distance

[16, 6] under Bregman’s general framework for convex optimization [5]. This framework is closely related to the later work of I. Csiszár and G. Tusnády [7] and generalizes much the earlier work [1, 19] and also [10, 4].

Our main computational focus in this paper is on the semidefinite programming problem corresponding to MAX CUT relaxation, i.e. with diagonal constraints. Another class of problems which allows basically the same complexity and convergence analysis corresponds to block-diagonal positive-definite constraints. The case of two complimentary block-diagonal constraints can be solved using Singular Value Decomposition, while our method in this case is a Riccati-like iteration. We believe that beyond two block-diagonal constraints that our method is the best currently available practical method for large problems.

Another application of our method is the efficient solution of *matrix completion problems* [13]. Here the objective is to find the maximum determinant positive matrix with specified values on and near to its diagonal. This corresponds to finding the maximum entropy normal density having prescribed marginal distributions.

Finally, we want to clarify one important historical point: our choice of a barrier $-\log(\det P)$ was motivated by Bregman’s seminal paper [5]. In a discrete setting this approach also has many merits, but lacks the crucial property *self-concordance* [21] property. But in the Gaussian case this barrier is (basically Bregman’s) is a sense optimal [21]. This celebrated self-concordance property has been studied in the context of Newton-like methods. We used it to analyze the convergence of our projection-like method.

2 The Basic Framework

Given a joint probability distribution $p(x, y)$ on a product space $X \times Y$, and another distribution $q(x)$ defined on X alone, we define a new distribution $p'(x, y) \triangleq p(y|x)q(x)$, which inherits as much as possible from p while forcing its x -marginal to match q . The focus of our paper is exposing certain properties and algorithms relating to this natural operation in which a joint distribution is *deflated* via conditionalization, and then *reinflated* by multiplication to exhibit some required marginal.

We denote by p_x the x -marginal of $p(x, y)$, i.e. $\int p(x, y)dy$. Similarly we denote by $p_{y|x}$ the function p/p_x , i.e. the conditional probability of y given x . Next, given distribution $q(x)$ on X and denoting it q_x we define the operator:

$$S_{x,q}(p) \triangleq p_{y|x} \cdot q_x \tag{1}$$

which transforms p into a new distribution whose x -marginal is exactly q .

Among all distributions having this marginal, $S_{x,q}(p)$ is special in as much as no other such distribution is closer to p under the Kullback-Leibler (KL) distance [16, 6] denoted $D(\|\cdot\|)$. That is, $S_{x,q}(p)$ is a *projector* with respect to KL-distance and the set of distributions with x -marginal q . This is formalized by:

Proposition 1 *For a given distribution $p(x, y)$ the minimum of $D(r\|p)$ over distributions r having x -marginal $q(x)$ is realized by $S_{x,q}(p)$.*

proof: We begin by establishing the following important *Pythagorean* KL-distance identity, which we will refer to again later in the paper:

$$D(q\|p) - D(q\|p_{y|x} \cdot q_x) = D(q_x\|p_x) \quad (2)$$

for joint distributions p and q . From the definition of KL-distance and the observation that $p(x, y) = p_{y|x}(x, y) \cdot p_x(x)$, the left side of this equality becomes:

$$\begin{aligned} & \int q(x, y) \log \frac{q(x, y)}{p(x, y)} dx dy - \\ & \int q(x, y) \log \frac{q(x, y)}{p_{y|x}(x, y) \cdot q_x(x)} dx dy = \\ & \int q(x, y) \log \frac{p_{y|x}(x, y) \cdot q_x(x)}{p_{y|x}(x, y) \cdot p_x(x)} dx dy = \\ & \int q_x(x) \log \frac{q_x(x)}{p_x(x)} dx = D(q_x\|p_x) \end{aligned}$$

Since by assumption $r_x = q(x)$ we may then write:

$$D(r\|p) = D(q(x)\|p_x) + D(r\|p_{y|x} \cdot q(x))$$

which is clearly minimized by choosing r to be $p_{y|x} \cdot q(x)$, i.e. $S_{x,q}(p)$, since KL-distance is zero when its two arguments are the same distribution. \square

Now $S_{x,q}(p)$ is defined in Eq-1 for product spaces $X \times Y$. But the definition easily extends to more general finite product spaces $X_1 \times X_2 \times \dots \times X_n$ by letting X stand for some subset of the dimensions $1, 2, \dots, n$, and Y stand for the others. The function q is then defined on the product space corresponding to the members of X .

More formally, let $K = i_1, i_2, \dots, i_k$ be a subset of the dimensions $1, 2, \dots, n$ and q be a function of $X_{i_1} \times X_{i_2} \times \dots, X_{i_k}$. Then the operator $S_{K,q}(p)$ transforms p into a new function whose K -marginal exactly matches q .

We view the pair (K, q) as a constraint on the behavior of a desired distribution. Then given a starting distribution p , and a collection $(K_1, q_1), (K_2, q_2), \dots, (K_m, q_m)$

of such constraints, a natural approach to the transformation of p into a distribution that satisfies all of them, is to iteratively apply the S operator — cycling through the collection of constraints. One cycle corresponds to the composition:

$$S_{K_m, q_m} \circ \dots \circ S_{K_2, q_2} \circ S_{K_1, q_1} \quad (3)$$

and many such cycles may be needed.

We remark that in a particular discrete probability setting this iteration corresponds to Sinkhorn balancing [23, 24]. That is, where $X_1 = \{1, 2, \dots, N\}$, $X_2 = \{1, 2, \dots, N\}$, $K_1 = \{1\}$, and $K_2 = \{2\}$. The distribution is then represented by an $N \times N$ matrix with nonnegative entries and the iteration corresponds to an alternating row-wise and column-wise renormalization procedure. In Sinkhorn's original paper q_1 and q_2 are uniform and simply scale down the entire matrix so that the sum of its entries is unity. There is a considerable literature on this subject, but it is not the focus of our paper.

In the Sinkhorn case, the iteration of Eq-3 gives rise to a constructive algorithm because the simple discrete distributions involved have the property that:

1. the S operator may be conveniently computed, and;
2. its result is again a member of the same class of distributions.

Normal densities represent another important class for which these properties hold and here too the approach of Eq-3 leads to a simple and concrete algorithm. Consider the zero-mean normal distribution on \mathbb{R}^n :

$$N_P(z) \triangleq \frac{\sqrt{\det P^{-1}}}{(2\pi)^{n/2}} \cdot e^{-\frac{1}{2}z^t P^{-1}z}$$

and assume that the n dimensions are partitioned into two sets X and X^\perp . Next, without loss of generality assume that X corresponds to dimensions $1, 2, \dots, r$ and that X^\perp corresponds to $r + 1, r + 2, \dots, n$. Then write P and P^{-1} in block form as:

$$P = \begin{pmatrix} X & X^\perp \\ P_{11} & P_{12} \\ P_{12}^t & P_{22} \end{pmatrix} \quad P^{-1} = \begin{pmatrix} X & X^\perp \\ \tilde{P}_{11} & \tilde{P}_{12} \\ \tilde{P}_{12}^t & \tilde{P}_{22} \end{pmatrix}$$

Now for any $Q \succ 0$ consider normal density N_Q on \mathbb{R}^r corresponding to subspace X . The marginal of N_P corresponding to subspace X is $N_{P_{11}}$. So $S_{X, N_Q}(N_P) = N_P N_Q / N_{P_{11}}$. Ignoring the leading constant and the $1/2$ term in the exponent the normal density corresponding to S may be written:

$$e^{-(x^t \tilde{P}_{11} x + 2x^t \tilde{P}_{12} x^\perp + x^\perp \tilde{P}_{22} x^\perp)} \cdot e^{-x^t Q^{-1} x} / e^{-x^t P_{11}^{-1} x}$$

from which it is clear that $S_{X, N_Q}(N_P) = N_A$ where

$$A^{-1} = \begin{pmatrix} X & X^\perp \\ \tilde{P}_{11} - (P_{11})^{-1} + Q^{-1} & \tilde{P}_{12} \\ \tilde{P}_{12}^t & \tilde{P}_{22} \end{pmatrix} \quad (4)$$

and for brevity in what follows we abuse our S notation and write:

$$S_{X, Q}(P) = A$$

It is important to notice that this operation is best carried out using the inverses of the matrices involved. That is, given P^{-1} one arrives at A^{-1} by simply adding $Q^{-1} - (P_{11})^{-1}$ to its upper left block.

The next proposition connects our probabilistic development involving normal densities, with the simpler notion of positive definite matrices and their determinants. It establishes that the Bregman distance [5, equation 1.4] arising from the determinant functional operating on positive definite matrices P_1, P_2 , is essentially the same as the Kullback-Leibler distance between the zero-mean normal densities arising from these matrices.

Proposition 2 Define $\psi(P) \triangleq -\log(\det P)$, for $n \times n$ matrices $P \succ 0$, and consider Bregman's distance

$$B(P_2, P_1) \triangleq \psi(P_1) - \psi(P_2) - \langle \nabla \psi|_{P=P_2}, P_1 - P_2 \rangle \quad (6)$$

then $D(N(P_2) \| N(P_1)) = 1/2 \cdot B_\psi(P_2, P_1)$, and ψ is a strictly convex function.

proof: It is well known that $D(N(P_2) \| N(P_1)) = 1/2(-\log(\det P_1) + \log(\det P_2) + \text{tr}(P_1 - P_2)P_2^{-1}) = 1/2(\text{tr}P_1P_2^{-1} - \log(\det P_1P_2^{-1}) - n)$. It is also well known that $\nabla \psi = -P^{-1}$. Next recall that the (Frobenius) inner product of two symmetric matrices is just the trace of their product [14]. So $B_\psi(P_2, P_1) = \text{tr}P_1P_2^{-1} - \log(\det P_1P_2^{-1}) - n$, whence $D(N(P_2) \| N(P_1)) = 1/2B(P_2, P_1)$. Now from our knowledge of normal densities and KL-distance we have that $D(N(P_2), N(P_1))$ is nonnegative, and positive iff $P_2 \neq P_1$ — from which it follows that B_ψ is positive except when its arguments are equal. Since B_ψ is continuous and differentiable its strict convexity is established, i.e., the function's graph lies strictly to one side of a tangent plane. We remark that convexity may also be proven by expressing B_ψ in the basis where $P_1P_2^{-1}$ is diagonal, noticing that all eigenvalues must be real and positive, and using the inequality $x - \log x - 1 \geq 0$. \square

In this light, our iterative approach is seen to be a Bregman projection method. That is, proposition 1 shows that the S operator minimizes KL distance between normal densities, but this is the same as projection under Bregman distance operating on positive definite matrices using the determinant functional.

So applying the operator $S_{X, Q}$ to a matrix $P \succ 0$ gives the matrix D that is the closest to P while agreeing with Q on subspace X . Given this observation, the natural algorithm for discovering a matrix that satisfies a system of such constraints, is to cycle over them using the S operator to forcibly impose each constraint in turn.

Our first theorem states that this natural algorithm works provided that a solution exists, and that it discovers the closest solution to the initial matrix chosen. Its proof requires the following observation which establishes that all members of a bounded Bregman neighborhood are themselves bounded above, and bounded strictly away from zero below.

Proposition 3 Let D be a positive definite matrix and consider a Bregman neighborhood $N_{D, r} \triangleq \{P | P \succ 0, B(P, D) \geq r\}$ of radius r about D . Then there exists ℓ, u such that $0 < \ell I < P < uI, \forall P \in N_{D, r}$.

proof: From the definition of Bregman distance we write $B(P, D) = \log \det P + \text{tr}DP^{-1} - (n + \log \det D)$. We are free to choose a basis in which P is diagonal since $B(P, D)$ is easily seen to be basis-invariant. This follows from direct calculation, or, more interestingly by reasoning that Bregman distance is essentially the same as KL distance, which is an entropic property of two normal densities that fairly clearly coordinate system independent. Then $B(P, D) \geq \log \bar{\lambda}(P) + C$ where $\bar{\lambda}(P)$ denotes the largest eigenvalue of P , and C is a constant depending on D . So $B(P, D) \rightarrow \infty$ as $\bar{\lambda}(P) \rightarrow \infty$. Next let $\mu > 0$ denote the smallest diagonal element of D . Then $B(P, D) \geq n \log \underline{\lambda}(P) + \mu / \underline{\lambda}(P)$. Here $B(P, D) \rightarrow \infty$ as $\underline{\lambda}(P) \rightarrow 0$, whence elements of any bounded Bregman neighborhood must have spectra bounded above, and bounded strictly away from zero below \square

Theorem 1 Consider a set $(X_i, Q_i), 0 \leq i < m$, of linear subspaces $X_i \subset \mathbb{R}^n$ and linear mapping constraints $0 < Q_i : X_i \rightarrow X_i$. If there exists $0 < D : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that $D_{X_i, X_i} = Q_i$, then for any initial matrix $P_0 \succ 0$, the algorithm that repeats the cyclic composition:

$$S_{X_{m-1}, Q_{m-1}} \circ \dots \circ S_{X_1, Q_1} \circ S_{X_0, Q_0}$$

starting with P_0 , converges to a limit $\bar{P} \succ 0$ such that:

1. \bar{P} satisfies the constraints, and;

2. $B(P_0, D) = B(P_0, \bar{P}) + B(\bar{P}, D)$, $\forall D \succ 0$ satisfying the constraints.

proof: We write the cyclic iteration as:

$$\begin{aligned} P_1 &= S_{X_0, Q_0}(P_0) \\ &\vdots \\ P_{n+1} &= S_{X_{[n]_m}, Q_{[n]_m}}(P_n) \end{aligned}$$

Then for any positive definite matrix $D : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that $D_{X_i, X_i} = Q_i$, $i = 0, \dots, m-1$, we have from equation 2 that:

$$\begin{aligned} 0 \leq D(P_1, D) &= D(P_0, D) - D([P_0]_{X_0, X_0}, Q_0) \\ &\vdots \\ 0 \leq D(P_{n+1}, D) &= D(P_n, D) - D([P_n]_{X_{[n]_m}, X_{[n]_m}}, Q_{[n]_m}) \end{aligned}$$

It is then clear that $0 \leq D(P_{n+1}, D) \leq D(P_n, D) \leq \dots \leq D(P_0, D) < \infty$, for all $n > 0$, and that:

$$\sum_{n=0}^{\infty} D([P_n]_{X_{[n]_m}, X_{[n]_m}}, Q_{[n]_m}) < \infty$$

Now the set $\{P_n\}$ of iterates has bounded Bregman distance to D , so by proposition 3 this set is bounded. Therefore a concentration point must exist. Since $\{P_n\}$ is bounded strictly away from zero, its closure, and therefore all concentration points must be positive definite.

Next observe that any concentration point must be stationary under S since this is a continuous mapping (in P). It then follows that there is a unique concentration point which we denote by \bar{P} . Moreover, its stationarity implies that it satisfies all constraints.

Now it is easy to see that:

$$B(\bar{P}, D) = B(P_0, D) - \sum_{n=0}^{\infty} B([P_n]_{X_{[n]_m}, X_{[n]_m}}, Q_{[n]_m})$$

Letting $D = \bar{P}$ this implies that $B(P_0, \bar{P}) = \sum_{n=0}^{\infty} B([P_n]_{X_{[n]_m}, X_{[n]_m}}, Q_{[n]_m})$. But then $B(P_0, D) = B(P_0, \bar{P}) + B(\bar{P}, D)$ \square

We remark that a similar theorem and proof exist for alternatives to cyclic iteration such as the *greedy* approach in which a subspace having maximal distance to the target distribution is selected at each iteration.

In the Bregman framework, the constraint set corresponds to an intersection of convex subspaces, and the iteration converges to the point in this intersection closest to the starting point. Elements of the intersection

(matrices satisfying all constraints) are stationary under the iteration.

A lot more can be said about the iteration of theorem 1. Later we discuss its rate of convergence, describe cases for which it converges in a finite number of iterations, prove that similar results hold for alternatives to the cyclic constraint ordering, and give a setting under which its computation is particularly fast.

When the intersection consists of more than one point Bregman observes that his iterative projection framework can do more than discover an arbitrary one — it can be guided, through appropriate selection of an initial value, to converge on the point that solves some minimization problem over the intersection region. In our case it can maximize the determinant, and we will show later that our algorithm may be used to rapidly compute the maximum determinant positive definite completion of a band matrix.

But our next step is to directly connect theorem 1 to certain problems of *semidefinite programming*. See [12, 11] for a review and connections to combinatorial optimization. In what follows we denote Frobenius matrix inner product by $A \bullet B$. The semidefinite programming problem we address is:

Given $C \succ 0$, find a matrix P that minimizes $C \bullet P$, subject to $P \succ 0$ and $P_{X_i, X_i} = Q_i$, $i = 0, \dots, m-1$.

and we will refer to it as CLS-SDP, standing for *constrained linear subspace semidefinite programming*. We observe that the semidefinite relaxation of MAX CUT given in [11] is an instance of CLS-SDP. The following proposition connects the distance minimization goal of theorem 1 to this problem.

Proposition 4 Given $C \succ 0$, minimizing $B(\epsilon C^{-1}, P)$ is the same as minimizing $C \bullet P - \epsilon \log \det P$

proof: Given $\epsilon > 0$, we have from Eq 6:

$$B(\epsilon C^{-1}, P) = \text{tr} P \frac{C}{\epsilon} - \log \det P - \log \det \frac{C}{\epsilon} - N$$

for any $P \succ 0$. This may be written:

$$\epsilon B(\epsilon C^{-1}, P) = \text{tr} PC - \epsilon \log \det P - \underbrace{\epsilon \log \det \frac{C}{\epsilon} - \epsilon N}_{\text{constant}}$$

where the last two terms are constant with respect to our minimization of P . It is then clear that minimizing $B(\epsilon C^{-1}, P)$ is the same as minimizing $\text{tr} PC - \epsilon \log \det P$. \square

As $\epsilon \rightarrow 0$ the role of $\log \det P$ diminishes and the focus is increasingly on $\text{tr} PC$, or $P \bullet C$. We remark that when

the constraints correspond to central minors of a single band matrix, starting points of the form $\frac{1}{\epsilon}(C+D)^{-1}$ can be shown to function in the same way where D is any nonnegative diagonal matrix.

Now given $\epsilon > 0$ denote by P_ϵ the unique solution to the minimization of $C \bullet P - \epsilon \log \det P$. Next let α denote the minimum value of $C \bullet P$ achieved for an instance of CLS-SDP. It is easy to see that if $\epsilon_1 > \epsilon_2$ then $C \bullet P_{\epsilon_1} \geq C \bullet P_{\epsilon_2} \geq \alpha$. That P_ϵ may be viewed as an approximate solution to the original problem is established by proposition 6, but we first establish a technical result needed for its proof.

Proposition 5 *Given C , $\det P_\epsilon$ is bounded above independent of $\epsilon \in (0, 1]$.*

proof: We will bound $\bar{\lambda}(P_\epsilon)$, the largest eigenvalue of P_ϵ . Choose some $\hat{P} \succ 0$ satisfying the constraints. Then $\text{tr}CP_\epsilon - \epsilon \log \det P_\epsilon \leq \text{tr}C\hat{P} - \epsilon \log \det \hat{P} \leq \text{tr}C\hat{P} + |\log \det \hat{P}|$ establishing an upper bound on $\text{tr}CP_\epsilon - \epsilon \log \det P_\epsilon$ independent of ϵ .

For any unit length vector u it is easy to show that $(u, Cu) > \underline{\lambda}(C)$, the smallest eigenvalue of C . It follows that $\underline{\lambda}(C)$ provides a lower bound on the diagonal elements of C that holds for any basis in which C is expressed. So choosing a basis that diagonalizes P_ϵ we have that $\text{tr}CP_\epsilon \geq \underline{\lambda}(C)\bar{\lambda}(P_\epsilon)$. It also follows that $\epsilon \log \det P_\epsilon \leq \epsilon N \log \bar{\lambda}(P_\epsilon)$. Combining these we establish the lower bound $\text{tr}CP_\epsilon - \epsilon \log \det P_\epsilon \geq \underline{\lambda}(C)\bar{\lambda}(P_\epsilon) - \epsilon N \log \bar{\lambda}(P_\epsilon) \geq \underline{\lambda}(C)\bar{\lambda}(P_\epsilon) - |N \log \bar{\lambda}(P_\epsilon)|$. Eventually, as increasing $\bar{\lambda}(P_\epsilon)$ are postulated, the lower bound exceeds the upper bound — establishing an upper bound on $\bar{\lambda}(P_\epsilon)$ independent of ϵ . \square

Proposition 6 *In the setting above we have: $0 < C \bullet P_\epsilon - \alpha \leq O(\epsilon \log \frac{1}{\epsilon})$.*

proof: Select some $\hat{P} \succ 0$ satisfying the constraints and choose u, ℓ such that $uI \succ \hat{P} \succ \ell I > 0$. Next let \bar{P} denote a solution to the CLS-SDP instance. That is, a matrix that satisfies the constraints such that $C \bullet \bar{P}$ achieves its minimal value α . Define $\hat{P}_\epsilon = (1 - \epsilon)\bar{P} + \epsilon\hat{P}$. Notice \hat{P}_ϵ satisfies the constraints, that $\hat{P}_\epsilon \succ 0$ provided $\epsilon > 0$, and that $\hat{P}_\epsilon \geq \epsilon \ell I$. Now:

$$\begin{aligned} & \text{tr}C\hat{P}_\epsilon - \epsilon \log \det \hat{P}_\epsilon \\ &= (1 - \epsilon)\text{tr}(C\bar{P}) + \epsilon \text{tr}C\hat{P} - \epsilon \log \det \hat{P}_\epsilon \\ &\leq (1 - \epsilon)\alpha + \epsilon u \|C\| + \epsilon N \log \frac{1}{\epsilon \ell} \\ &= \alpha + \epsilon(u \|C\| - \alpha) + \epsilon N \log \frac{1}{\epsilon \ell} \\ &= \alpha + O(\epsilon \log \frac{1}{\epsilon}) \end{aligned}$$

Now by definition $\text{tr}CP_\epsilon - \epsilon \log \det P_\epsilon \leq \text{tr}C\hat{P}_\epsilon - \epsilon \log \det \hat{P}_\epsilon \leq \alpha + O(\epsilon \log \frac{1}{\epsilon})$. From proposition 5 we have that $\log \det P_\epsilon = O(1)$ so $\text{tr}CP_\epsilon - \alpha \leq O(\epsilon \log \frac{1}{\epsilon}) + \epsilon O(1) = O(\epsilon \log \frac{1}{\epsilon})$ \square

3 The Band Matrix Setting

The “band matrix” setting provides an elegant application of proposition 4 and theorem 1. That is, $n \times n$ matrices with nonzero entries confined to distance m from the diagonal. Propositions 8 gives a key result for this setting and the following general observation is key to both of them. It tells us that when a marginal constraint that *straddles* the two components of a product distribution is forced by applying S , then the result exhibits marginals that agree with the component distributions provided that the constraint agrees with a given component on the region of overlap.

Proposition 7 *Let $p(x_1, x_2, x_3, x_4)$ be a probability function that may be written as a product $f(x_1, x_2)g(x_3, x_4)$. Next let $q(x_2, x_3)$ be a probability function thought of as a constraint. Then the X_1, X_2 marginal of $S_{\{X_2, X_3\}, q}(p)$ remains $f(x_1, x_2)$ provided that $q(x_2) = f(x_2)$. That is, provided that q is consistent with f on the overlapping variable. Similarly the X_3, X_4 marginal of $S_{\{X_2, X_3\}, q}(p)$ remains $g(x_3, x_4)$ provided that $q(x_3) = g(x_3)$.*

proof: We write:

$$S_{\{X_2, X_3\}, q}(p) = q(x_2, x_3) \frac{f(x_1, x_2)g(x_3, x_4)}{f(x_2)g(x_3)}$$

To evaluate its X_1, X_2 marginal we first integrate over X_4 yielding $q(x_2, x_3) f(x_1, x_2)/f(x_2)$, and then integrate over X_3 giving $q(x_2)f(x_1, x_2)/f(x_2)$. Now if $q(x_2) = f(x_2)$ then the result is $f(x_1, x_2)$ as desired. The same argument applies to the X_3, X_4 marginal case. \square

Our first band-matrix result applies the development above to solve the “band matrix completion problem.” [13]. It is fascinating that we can reason about this rather pure matrix problem and so easily solve it using a probabilistic outlook and the language of normal densities.

Proposition 8 *Given a band-structured positive definite matrix C , a closed-form formula exists for its unique maximum determinant completion, and this formula corresponds to exactly $n - m + 1$ iterations of the deflation-inflation computation of theorem 1.*

proof: We first argue that the process terminates in $n - m + 1$ steps. The sequence of computation is depicted in figure 1. Notice that since the starting matrix is diagonal, and each application of S only affects a central minor, the final result must also have band structure. By definition the first application of S enforces the constraint corresponding to the first central minor of C .

After k applications we consider the space to be a product of two subspaces. The first corresponding to diagonal elements $1, \dots, m + k - 1$, and the second to elements $m + k, \dots, n$. Since the later portion of the matrix is diagonal, it corresponds to the product of unit variance independent normal densities, and the product-form requirement of proposition 7 is satisfied.

The constraint corresponding to the following application ($k + 1$) of S is consistent with the constraints enforced earlier because the overlap of its corresponding central minor of C with the $1, \dots, m + k - 1$ subspace is a submatrix of the central minors already enforced. Thus, the $q(x_2) = f(x_2)$ of the proposition is satisfied. It then follows that application $k + 1$ will not disturb the marginals already enforced — so after $n + m - 1$ applications all constraints are satisfied and we denote the result \bar{P} .

Now by theorem 1 \bar{P} minimizes the Bregman distance $I \bullet P - \log \det P$ subject to the constraints given by the central minors of C . But the diagonal of any P satisfying them must match the diagonal of C , so the $I \bullet P$ term is constant in the minimization, whence \bar{P} maximizes $\log \det P$. \square

A naive implementation of the computation of proposition 8 requires $O(n^4)$ time. But if A is an $n \times n$ m -band matrix, recall that the linear system $Ax = b$ can be solved in $O(n)$ operations. If A is nonsingular, its inverse can then be constructed column by column, and any constant-sized sub-block is obtained in $O(n)$ time. So regarding m as fixed, it is then clear that the computation of proposition 8 is reduced to $O(n^2)$ time. That is, \bar{P}^{-1} (the inverse of the maximum-determinant completion) may be found in $O(n^2)$ time. Since \bar{P}^{-1} is band-structured, it may be inverted in $O(n^2)$ time to yield \bar{P} , so that the entire problem has this complexity.

The concrete example above builds intuition, but by instead framing the problem in more abstract terms we are quickly led to an $O(n)$ time highly general solution.

Given marginals $q^i(x_i, \dots, x_{i+d-1})$ on $[1, \dots, d], [2, \dots, d + 1], \dots, [n - d + 1, n]$, where we assume without loss of generality that $d = 2m$ and $n = 2m'$, we can group variables as $[x_1, \dots, x_k]$ where $x_1 = [1, \dots, d/2], x_2 = [d/2 + 1, d]$, and so on. One may think of this as a *block tridiagonal* representation of the system of marginals.




Figure 1: The inverse of the maximum determinant completion of a positive definite band matrix is constructed by starting with the identity matrix and applying the S operator to successive central minors. After $n - m + 1$ applications the process is complete, and may be interpreted as a closed form expression for the final result.

Starting with any product distribution $f_0(x_1, \dots, x_k) = s_1(x_1) \cdot \dots \cdot s_k(x_k)$ we perform the following Bregman-Sinkhorn iteration:

$$f_{i+1}(x_1, \dots, x_k) \rightarrow \frac{f_i(x_1, \dots, x_k) q^{(i)}(x_i, x_{i+1})}{\int f_i(x_1, \dots, x_k) dx_1 \dots dx_k}$$

Then the resulting maximum entropy distribution is by induction:

$$q^1(x_1, x_2) \cdot q^2(x_2 | x_1) \cdot \dots \cdot q^{n-d+1}(x_k | x_{k-1})$$

In the Gaussian case we have $q^i = N(Q^i)$ where:

$$Q^i = \begin{pmatrix} P_i & B_i \\ B_i^t & R_i \end{pmatrix} (Q^i)^{-1} = \begin{pmatrix} D_i & T_i \\ T_i^t & L_i \end{pmatrix}$$

Then the maximal determinant completion P has block tridiagonal inverse structure with diagonal blocks A_1, \dots, A_k and off-diagonal blocks T_1, \dots, T_{k-1} — where $A_1 = D_1, A_2 = L_1 + D_2 - P_2, A_3 = L_2 + D_3 - P_2, \dots, A_k = L_{k-1}$.

The result may be computed in $O(nd^2)$ time, or only $O(d^3)$ parallel time with k processors.

4 Local Analysis

The local convergence properties of our positive definite matrix iteration, and those of other instances of deflation-inflation, may be analyzed using the following

Lemma 1 Consider a $f : B \mapsto B$ where B is an open set of $\mathbb{R}^n \oplus \mathbb{R}^m$ containing the origin. Further assume that $f(x, 0) = (x, 0)$, and that for a particular $\bar{z} = (\bar{x}, 0)$, f is analytic in a neighborhood about it and the Jacobian of f has the following structure:

$$\frac{\partial f}{\partial z} \Big|_{\bar{z}} = \begin{pmatrix} \mathbb{R}^n & \mathbb{R}^m \\ I & 0 \\ 0 & A \end{pmatrix}$$

where $\|A\| = \alpha < 1$. Finally, assume that for some initial point z_0 , the sequence of iterates defined by $z_{n+1} = f(z_n)$ converges to $\bar{z} = (\bar{x}, 0)$. Then the rate of convergence is linear.

proof: Convergence implies that for sufficiently large n , $z_n = (x_n, y_n)$ where $\|y_n\| \ll 1$ and $\|x_n - \bar{x}\| \ll 1$ and f is analytic at $(x_n, 0)$. Thus $z_{n+1} = f((x_n, 0) + (0, y_n)) = (x_n, 0) + J_n \cdot (0, y_n) + O(\|y_n\|^2)$, where $J_n \triangleq \frac{\partial f}{\partial z} \Big|_{(x_n, 0)}$. Observe that the condition $f(x, 0) = (x, 0)$ implies

$$J_n = \begin{pmatrix} I & B_n \\ 0 & A_n \end{pmatrix}$$

where $A_n \rightarrow A$ and $B_n \rightarrow 0$ as $n \rightarrow \infty$. Next observe that $y_{n+1} = A_n y_n + O(\|y_n\|^2)$ so that $\exists \epsilon > 0$ such that $\|y_{n+1}\| \leq (\alpha + \epsilon)\|y_n\|$ for all n sufficiently large, where $(\alpha + \epsilon) < 1$. We also have $x_{n+1} = x_n + B_n y_n + O(\|y_n\|^2)$ So $x_{n-1} - x_n = O(\|y_n\|)$ establishing overall linear convergence. \square

Corollaries to these lemmas establish linear convergence for our normal-density inspired positive definite case, and provide a concise proof for a rather general form of Sinkhorn balancing. For a general Sinkhorn balancing we prove that the rate of convergence is linear iff the initial zero pattern is equal to the zero pattern of the limit. The statement and proofs of these results will be presented in the full paper.

5 Global Analysis

Beyond local rate of convergence we are interested in a bound on the total number of operations required to produce an approximate solution. Our analysis applies to a form of the algorithm in which the next diagonal position to update is chosen *greedily*, rather than in a simple cyclic fashion. That is, we update the diagonal element that generates the largest immediate reduction in KL distance. A similar analysis is possible for the somewhat more practical ordering based on KL distance for each marginal, i.e. for each diagonal element.

The analysis requires a number of technical lemmas and is rather involved. We sketch its development here leaving the details for the full paper. The main theorem we establish is:

Theorem 2 Deflation-Inflation in its greedy form with appropriate scheduling of ϵ reductions, requires $O(1/\epsilon \log 1/\epsilon)$ diagonal updates to produce a solution that leads to cuts that are on expectation within a factor of $(1 + \epsilon)$ of the value arising from the Goemans-Williamson analysis.

As the algorithm runs, the diagonal of C^{-1} approaches ϵI . It is first shown that the error in diagonal elements must be no more than $\epsilon^{3/2}$ to achieve the desired $(1 + \epsilon)$ approximation. This results from the observation that it is enough to continue until the KL distance to the optimal solution is $O(\epsilon^2)$.

Next, the initial KL distance is easily shown to be $O(1/\epsilon)$. We denote by μ the maximum diagonal error we will accept. That is, the algorithm will terminate once the maximum diagonal error less than μ . Then from proposition 1 it follows that every diagonal update, up until we stop, must generate at least $O(\mu^2)$ improvement in KL distance. So $O(1/\epsilon\mu^2)$ steps are required. Letting $\mu = \epsilon^{3/2}$ we arrive at an easy $O(1/\epsilon^4)$ naive bound.

We remark that this naive analysis may be adapted to the Sinkhorn case.

Fortunately, much stronger bounds may be developed for our normal density setting. The first step consists of letting $\mu = \sqrt{\epsilon}$ and showing that in $O(1/\epsilon^2)$ steps the solution comes within a unit KL ball.

Once within this ball it may be shown that only $O(1/\epsilon \log 1/\epsilon)$ steps are required to come within any fixed power of ϵ distance from the optimal solution – in particular within $\epsilon^{3/2}$. This improves our bound to $O(1/\epsilon^2)$, but additional improvement is possible.

The final idea consists of appropriate ϵ scheduling. This improves the first phase of the algorithm where the objective is to come within unit distance. We employ a series of shrinking ϵ values where each is a constant factor smaller than the previous one. This constant is of the form $(1 - c/n)$ for some constant c . For each value ϵ_i we perform $O(1/\epsilon_i)$ diagonal operations bringing us close enough to P_{ϵ_i} so that the ultimate effect is that after $O(1/\epsilon \log 1/\epsilon)$ steps we come within unit distance of P_ϵ as required. The second phase of the algorithm then runs, completing the solution.

We remark that while we do not use the interior point method in the typical fashion as a *black box*, several aspects of our analysis borrow from ideas at the heart of it's analysis. Also, while our development follows an entirely different probabilistic line, it is interesting to

note that we arrive at the same *barrier* function, which appears there.

To give a more detailed view of the technical analysis we state main lemmas here without proof.

Lemma 2 $\lambda_{\min}(P_\epsilon) \geq O(\epsilon/n)$

Lemma 3 $\lambda_{\min}(A \circ B) \geq 1 - (1 - \lambda_{\min}(A))(1 - \lambda_{\min}(B))$, where $A \circ B$ denotes the Schur (component-wise) product of positive definite matrices A and B , each with unit diagonal.

Lemma 4 Define $P(\lambda) = (P_\epsilon^{-1} + \text{Diag}(\lambda_i))^{-1} > 0$ and $f(\lambda) = ((1 - P(\lambda)(1, 1), \dots, (1 - P(\lambda)(n, n)))$, then for all $\epsilon > 0$:

1. If $D(P(\lambda)\|P_\epsilon) > 1$ then $\|f(\lambda)\|^2 \geq O(\epsilon/n)$
2. If $D(P(\lambda)\|P_\epsilon) \leq 1$ then $\|f(\lambda)\|^2 \geq O(\epsilon/n)D(P(\lambda)\|P_\epsilon)$

Lemma 5 For any $n \times n$ $P > 0$ with unit diagonal, then $\langle (P \circ P)^{-1}e, e \rangle \geq n$, where $e = (1, \dots, 1)$.

As remarked above, essentially the same *naive* analysis applies to “greedy” Sinkhorn balancing. Suppose A_0 is a row-normalized square incidence matrix of a bipartite graph G .

On each step of our algorithm in its greedy form we choose that row or column having sum of elements S such that $|S - 1|$ is maximal, and divide all entries in it by S .

It follows directly from [5] that this algorithm converges if and only if G has a perfect matching.

Using Hall’s theorem it is easy to prove (and well known) that if at some iteration $|S - 1| < 1/(2n + 1)$ then a perfect matching exists.

Now suppose that B is any double-stochastic matrix with:

$$B(A_0\|B) = \sum b_{i,j} \log \frac{b_{i,j}}{A_0(i,j)} + A_0(i,j) - b_{i,j} < \infty$$

then from proposition 1, adapted to this distance, it follows that:

$$B(A_{n+1}\|B) = B(A_n\|B) - (S - 1 - \log S)$$

Thus the number of iterations required to reach the condition $|S - 1| < 1/(2n + 1)$ is at most $C \cdot B(A_0\|B) \cdot n^2$, since $B(A_n\|B) \geq 0$ for all $n \geq 1$, where C is some small universal constant.

Next, if B is a permutational matrix of a perfect matching for G , then $B(A_0\|B) = \sum(\log d_i)$, where d_i denotes the degree of vertex i .

We have therefore produced a balancing algorithm that checks the existence of a perfect matching requiring $O(n^4 \log m/n)$ time or $O(n^3 \log m/n)$ parallel steps given $O(n)$ processors. Here m denotes the number of edges in G .

6 Concluding Remarks

Our Deflation-Inflation approach can not address semidefinite programming problems in general form but we remark that our CLS-SDP variant includes problems that can not be formulated as conventional instances of semidefinite programming without the introduction of a large number of constraints. Namely, instances that involve constraints that are expressed in terms of transformed linear subspaces. We also emphasize that our method is particularly attractive for specially structured matrices, and in particular for band matrices.

In our probabalistic framework we use normal densities, which may be thought of as capturing second-order relationships between vertices – namely edges. We speculate that different probability functions might lead to interesting results. Other members of the exponential family might be considered, or higher-order generalizations of normal densities.

It is clear that any global analysis of projection-based convex optimization methods (such as EM) requires bounds on the relationship between global KL distance and marginal distances. During our analysis we remark that machinery from [21] was borrowed to allow us to establish some bounds of this sort. These *entropic inequalities* may be of independent interest.

Finally we remark that there is tension between work invested in generating random cuts, and work aimed at solving the semidefinite program for ever smaller ϵ values. Another interesting area for future work is to focus on the variance of the cuts generated in order to better understand this tradeoff.

References

- [1] AGMON, S. The relaxation method for linear inequalities. *Canadian Journal of Mathematics* 6, 3 (1954), 382–392.
- [2] ALIZADEH, F. Interior point methods in semidefinite programming with applications to combinatorial optimization. *SIAM Journal on Optimization* 5, 1 (1995), 13–51.
- [3] BENSON, S. J., YE, Y., AND ZHANG, X. Solving large-scale sparse semidefinite programs for combinatorial optimization. Manuscript, October 1997.
- [4] BREGMAN, L. M. Finding the common point of convex sets by the method of successive projections. *Dokl. Akad. Nauk SSSR* 162 (1965), 487–490.
- [5] BREGMAN, L. M. The relaxation method of finding the common point of convex sets and its application

- to the solution of problems in convex programming. *Zh. vĕchisl. Mat. mat. Fiz.* 7, 3 (1967).
- [6] COVER, T. M., AND THOMAS, J. A. *Elements of Information Theory*. Wiley, 1991.
- [7] CSISZÁR, I., AND TUSNÁDY, G. Information geometry and alternating minimization procedures. *Statistics & Decisions Supplement Issue No. 1* (1984), 205–237. R. Oldenbourg Verlag, München 1984.
- [8] DAHLQUIST, G., AND ÅKE BJÖRCK. *Numerical Methods*, ned anderson (translator) ed. Prentice-Hall, 1974.
- [9] DEMPSTER, A. P., LAIRD, N. M., AND RUBIN, D. B. Maximum-likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Society Ser. B (methodological)* 39 (1977), 1–38.
- [10] EREMIN, I. I. A generalization of the motzkin - agmon relaxational method. *Usp. Mat. Nauk* 20 (1965), 183–187.
- [11] GOEMANS, M. X. Semidefinite programming in combinatorial optimization. *Mathematical Programming* 79 (1997), 143–161.
- [12] GOEMANS, M. X., AND WILLIAMSON, D. P. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM* 42 (1995), 1115–1145.
- [13] GOHBERG, I., KAASHOEK, M. A., AND CAN SCHA- GEN, F. *Partially Specified Matrices and Operators: Classification, Completion, Applications*, vol. 79 of *Operator Theory: Advances and Applications*. Birkhäuser Verlag, Basel, 1995.
- [14] HORN, R. A., AND JOHNSON, C. R. *Matrix Analysis*. Cambridge University Press, 1985.
- [15] KLEIN, P., AND LU, H.-I. Efficient approximation algorithms for semidefinite programs arising from max cut and coloring. In *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing (STOC'96)* (1996), pp. 338–347.
- [16] KULLBACK, S., AND LEIBLER, R. A. On information and sufficiency. *Annals of Mathematical Statistics* 22 (1951), 79–86.
- [17] LOVÁSZ, L. On the shannon capacity of a graph. *IEEE Transactions on Information Theory* 25 (1979), 1–7.
- [18] LOVÁSZ, L., AND SCHRIJVER, A. Cones of matrices and setfunctions, and 0–1 optimization. *SIAM Journal on Optimization* 1 (1991), 166–190.
- [19] MOTZKIN, T. S., AND SCHOENBERG, I. I. The relaxation method for linear inequalities. *Canadian Journal of Mathematics* 6, 3 (1954), 393–404.
- [20] N. LINIAL, A. S., AND WIDGERSON, A. (private communication). 1998.
- [21] NESTEROV, Y., AND NEMIROVSKII, A. *Interior Point Polynomial Methods in Convex Programming: Theory and Applications*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1994.
- [22] POLJAK, S., AND TUZA, Z. Maximum cuts and largest bipartite subgraphs. In *Combinatorial Optimization*, W. Cook, L. Lovász, and P. Seymour, Eds., DIMACS Series in Discrete Mathematics and Computer Science. AMS, Providence, RI, 1995.
- [23] SINKHORN, R. A relationship between arbitrary positive matrices and doubly stochastic matrices. *Ann. Mathematical Statistics* 35 (1964), 876–879.
- [24] SOULES, G. W. The rate of convergence of sinkhorn balancing. *Linear Algebra and its Applications* 150 (1991), 3–40.