

Clustering Partially Observed Graphs via Convex Optimization

Yudong Chen

Ali Jalali

Sujay Sanghavi

*Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, TX 78712, USA*

YDCHE@UTEXAS.EDU

ALI@MAIL.UTEXAS.EDU

SANGHAVI@MAIL.UTEXAS.EDU

Huan Xu

*Department of Mechanical Engineering
National University of Singapore
Singapore 117575, SINGAPORE*

MPEXUH@NUS.EDU.SG

Editor: Marina Meila

Abstract

This paper considers the problem of clustering a partially observed unweighted graph—i.e., one where for some node pairs we know there is an edge between them, for some others we know there is no edge, and for the remaining we do not know whether or not there is an edge. We want to organize the nodes into disjoint clusters so that there is relatively dense (observed) connectivity within clusters, and sparse across clusters.

We take a novel yet natural approach to this problem, by focusing on finding the clustering that minimizes the number of “disagreements”—i.e., the sum of the number of (observed) missing edges within clusters, and (observed) present edges across clusters. Our algorithm uses convex optimization; its basis is a reduction of disagreement minimization to the problem of recovering an (unknown) low-rank matrix and an (unknown) sparse matrix from their partially observed sum. We evaluate the performance of our algorithm on the classical Planted Partition/Stochastic Block Model. Our main theorem provides sufficient conditions for the success of our algorithm as a function of the minimum cluster size, edge density and observation probability; in particular, the results characterize the tradeoff between the observation probability and the edge density gap. When there are a constant number of clusters of equal size, our results are optimal up to logarithmic factors.

Keywords: graph clustering, convex optimization, sparse and low-rank decomposition

1. Introduction

This paper is about the following task: given partial observation of an undirected unweighted graph, partition the nodes into disjoint clusters so that there are dense connections within clusters, and sparse connections across clusters. By partial observation, we mean that for some node pairs we know if there is an edge or not, and for the other node pairs we do not know—these pairs are *unobserved*. This problem arises in several fields across science and engineering. For example, in sponsored search, each cluster is a submarket that represents a specific group of advertisers that do most of their spending on a group of query phrases—see e.g., Yahoo!-Inc (2009) for such a project at Yahoo. In VLSI and design automation, it is useful in minimizing signaling between components (Kernighan and Lin, 1970). In social networks, clusters may represent groups of people with

similar interest or background; finding clusters enables better recommendations and link prediction (Mishra et al., 2007). In the analysis of document databases, clustering the citation graph is often an essential and informative first step (Ester et al., 1995). In this paper, we will focus not on specific application domains, but rather on the basic graph clustering problem itself.

Partially observed graphs appear in many applications. For example, in online social networks like Facebook, we observe an edge/no edge between two users when they accept each other as a friend or explicitly decline a friendship suggestion. For the other user pairs, however, we simply have no friendship information between them, which are thus unobserved. More generally, we have partial observations whenever obtaining similarity data is difficult or expensive (e.g., because it requires human participation). In these applications, it is often the case that *most* pairs are unobserved, which is the regime we are particularly interested in.

As with any clustering problem, we need a precise mathematical definition of the clustering criterion with potentially a guaranteed performance. There is relatively few existing results with provable performance guarantees for graph clustering with partially observed node pairs. Many existing approaches to clustering fully observed graphs either require an additional input (e.g., the number of clusters k required for spectral or k -means clustering methods), or do not guarantee the performance of the clustering. We review existing related work in Section 1.2.

1.1 Our Approach

We focus on a natural formulation, one that does not require any other extraneous input besides the graph itself. It is based on minimizing *disagreements*, which we now define. Consider any candidate clustering; this will have (a) observed node pairs that are in different clusters, but have an edge between them, and (b) observed node pairs that are in the same cluster, but do not have an edge between them. The total number of node pairs of types (a) and (b) is the number of disagreements between the clustering and the given graph. We focus on the problem of finding the optimal clustering—one that minimizes the number of disagreements. Note that we do *not* pre-specify the number of clusters. For the special case of fully observed graphs, this formulation is exactly the same as the problem of *correlation clustering*, first proposed by Bansal et al. (2002). They show that exact minimization of the above objective is NP-hard in the worst case—we survey and compare with this and other related work in Section 1.2. As we will see, our approach and results are different.

We aim to achieve the combinatorial disagreement minimization objective using matrix splitting via convex optimization. In particular, as we show in Section 2 below, one can represent the adjacency matrix of the given graph as the sum of an unknown low-rank matrix (corresponding to “ideal” clusters) and a sparse matrix (corresponding to disagreements from this “ideal” in the given graph). Our algorithm either returns a clustering, which is guaranteed to be disagreement minimizing, or returns a “failure”—it never returns a sub-optimal clustering. For our main analytical result, we evaluate our algorithm’s performance on the natural and classical *planted partition/stochastic block model* with partial observations. Our analysis provides stronger guarantees than are current results on general matrix splitting (Candès et al., 2011; Hsu et al., 2011; Li, 2013; Chen et al., 2013). The algorithm, model and results are given in Section 2. We prove our theoretical results in Section 3 and provide empirical results in Section 4.

1.2 Related Work

Our problem can be interpreted in the general clustering context as one in which the presence of an edge between two points indicates a “similarity”, and the lack of an edge means “no similarity”. The general field of clustering is of course vast, and a detailed survey of all methods therein is beyond our scope here. We focus instead on the three sets of papers most relevant to the problem here: the work on correlation clustering, on the planted partition/stochastic block model, and on graph clustering with partial observations.

1.2.1 CORRELATION CLUSTERING

As mentioned, for a completely observed graph, our problem is mathematically precisely the same as correlation clustering formulated in Bansal et al. (2002); in particular a “+” in correlation clustering corresponds to an edge in the graph, a “-” to the lack of an edge, and disagreements are defined in the same way. Thus, this paper can equivalently be considered as an algorithm, and guarantees, for *correlation clustering under partial observations*. Since correlation clustering is NP-hard, there has been much work on devising alternative approximation algorithms (Bansal et al., 2002; Emmanuel and Fiat, 2003). Approximations using convex optimization, including LP relaxation (Charikar et al., 2003; Demaine and Immorlica, 2003; Demaine et al., 2006) and SDP relaxation (Swamy, 2004; Mathieu and Schudy, 2010), possibly followed by rounding, have also been developed. We emphasize that we use a different convex relaxation, and we focus on understanding when our convex program yields an optimal clustering without further rounding.

We note that Mathieu and Schudy (2010) use a convex formulation with constraints enforcing positive semi-definiteness, triangle inequalities and fixed diagonal entries. For the fully observed case, their relaxation is at least as tight as ours, and since they add more constraints, it is possible that there are instances where their convex program works and ours does not. However, this seems hard to prove/disprove. Indeed, in the full observation setting they consider, their exact recovery guarantee is no better than ours. Moreover, as we argue in the next section, our guarantees are order-wise optimal in some important regimes and thus cannot be improved even with a tighter relaxation. Practically, our method is faster since, to the best of our knowledge, there is no low-complexity algorithm to deal with the $\Theta(n^3)$ triangle inequality constraints required by Mathieu and Schudy (2010). This means that our method can handle large graphs while their result is practically restricted to small ones (~ 100 nodes). In summary, their approach has higher computational complexity, and does not provide significant and characterizable performance gain in terms of exact cluster recovery.

1.2.2 PLANTED PARTITION MODEL

The planted partition model, also known as the stochastic block-model (Condon and Karp, 2001; Holland et al., 1983), assumes that the graph is generated with in-cluster edge probability p and inter-cluster edge probability q (where $p > q$) and fully observed. The goal is to recover the latent cluster structure. A class of this model with $\tau \triangleq \max\{1 - p, q\} < \frac{1}{2}$ is often used as benchmark for *average case* performance for correlation clustering (see, e.g., Mathieu and Schudy, 2010). Our theoretical results are applicable to this model and thus directly comparable with existing work in this area. A detailed comparison is provided in Table 1. For fully observed graphs, our result matches the previous best bounds in both the minimum cluster size and the difference between in-cluster/inter-cluster densities. We would like to point out that nuclear norm minimization has

been used to solve the closely related planted clique problem (Alon et al., 1998; Ames and Vavasis, 2011).

Paper	Cluster size K	Density difference $(1 - 2\tau)$
Boppana (1987)	$n/2$	$\tilde{\Omega}(\frac{1}{\sqrt{n}})$
Jerrum and Sorkin (1998)	$n/2$	$\tilde{\Omega}(\frac{1}{n^{1/6-\epsilon}})$
Condon and Karp (2001)	$\tilde{\Omega}(n)$	$\tilde{\Omega}(\frac{1}{n^{1/2-\epsilon}})$
Carson and Impagliazzo (2001)	$n/2$	$\tilde{\Omega}(\frac{1}{\sqrt{n}})$
Feige and Kilian (2001)	$n/2$	$\tilde{\Omega}(\frac{1}{\sqrt{n}})$
McSherry (2001)	$\tilde{\Omega}(n^{2/3})$	$\tilde{\Omega}(\sqrt{\frac{n^2}{K^3}})$
Bollobás and Scott (2004)	$\tilde{\Omega}(n)$	$\tilde{\Omega}(\sqrt{\frac{1}{n}})$
Giesen and Mitsche (2005)	$\tilde{\Omega}(\sqrt{n})$	$\tilde{\Omega}(\frac{\sqrt{n}}{K})$
Shamir and Tsur (2007)	$\tilde{\Omega}(\sqrt{n})$	$\tilde{\Omega}(\frac{\sqrt{n}}{K})$
Mathieu and Schudy (2010)	$\tilde{\Omega}(\sqrt{n})$	$\tilde{\Omega}(1)$
Rohe et al. (2011)	$\tilde{\Omega}(n^{3/4})$	$\tilde{\Omega}(\frac{n^{3/4}}{K})$
Oymak and Hassibi (2011)	$\tilde{\Omega}(\sqrt{n})$	$\tilde{\Omega}(\frac{\sqrt{n}}{K})$
Chaudhuri et al. (2012)	$\tilde{\Omega}(\sqrt{n})$	$\tilde{\Omega}(\frac{\sqrt{n}}{K})$
This paper	$\tilde{\Omega}(\sqrt{n})$	$\tilde{\Omega}(\frac{\sqrt{n}}{K})$

Table 1: *Comparison with literature.* This table shows the lower-bound requirements on the minimum cluster size K and the density difference $p - q = 1 - 2\tau$ that existing literature needs for exact recovery of the planted partitions, when the graph is fully observed and $\tau \triangleq \max\{1 - p, q\} = \Theta(1)$. Some of the results in the table only guarantee recovering the membership of most, instead of all, nodes. To compare with these results, we use the soft- Ω notation $\tilde{\Omega}(\cdot)$, which hides the logarithmic factors that are necessary for recovering all nodes, which is the goal of this paper.

1.2.3 PARTIALLY OBSERVED GRAPHS

The previous work listed in Table 1, except Oymak and Hassibi (2011), does not handle partial observations directly. One natural way to proceed is to impute the missing observations with no-edge, or random edges with symmetric probabilities, and then apply any of the results in Table 1. This approach, however, leads to sub-optimal results. Indeed, this is done explicitly by Oymak and Hassibi (2011). They require the probability of observation p_0 to satisfy $p_0 \gtrsim \frac{\sqrt{K_{\min}}}{n}$, where n is the number of nodes and K_{\min} is the minimum cluster size; in contrast, our approach only needs $p_0 \gtrsim \frac{n}{K_{\min}^2}$ (both right hand sides have to be less than 1, requiring $K_{\min} \gtrsim \sqrt{n}$, so the right hand side of our condition is order-wise smaller and thus less restrictive.) Shamir and Tishby (2011) deal with partial observations directly and shows that $p_0 \gtrsim \frac{1}{n}$ suffices for recovering two clusters of size $\Omega(n)$. Our result is applicable to much smaller clusters of size $\tilde{\Omega}(\sqrt{n})$. In addition, a nice feature of

our result is that it explicitly characterizes the tradeoffs between the three relevant parameters: p_0 , τ , and K_{\min} ; theoretical result like this is not available in previous work.

There exists other work that considers partial observations, but under rather different settings. For example, Balcan and Gupta (2010), Voevodski et al. (2010) and Krishnamurthy et al. (2012) consider the clustering problem where one samples the rows/columns of the adjacency matrix rather than its entries. Hunter and Strohmer (2010) consider partial observations in the features rather than in the similarity graph. Eriksson et al. (2011) show that $\tilde{\Omega}(n)$ actively selected pairwise similarities are sufficient for recovering a hierarchical clustering structure. Their results seem to rely on the hierarchical structure. When disagreements are present, the first split of the cluster tree can recover clusters of size $\Omega(n)$; our results allow smaller clusters. Moreover, they require active control over the observation process, while we assume random observations.

2. Main Results

Our setup for the graph clustering problem is as follows. We are given a partially observed graph of n nodes, whose adjacency matrix is $\mathbf{A} \in \mathbb{R}^{n \times n}$, which has $a_{i,j} = 1$ if there is an edge between nodes i and j , $a_{i,j} = 0$ if there is no edge, and $a_{i,j} = \text{"?"}$ if we do not know. (Here we follow the convention that $a_{i,i} = 0$ for all i .) Let $\Omega_{\text{obs}} \triangleq \{(i, j) : a_{i,j} \neq \text{"?"}\}$ be the set of observed node pairs. The goal is to find the optimal clustering, i.e., the one that has the minimum number of disagreements (defined in Section 1.1) in Ω_{obs} .

In the rest of this section, we present our algorithm for the above task and analyze its performance under the planted partition model with partial observations. We also study the optimality of the performance of our algorithm by deriving a necessary condition for any algorithm to succeed.

2.1 Algorithm

Our algorithm is based on convex optimization, and either (a) outputs a clustering that is guaranteed to be the one that minimizes the number of observed disagreements, or (b) declares “failure”. In particular, it never produces a suboptimal clustering.¹ We now briefly present the main idea and then describe the algorithm.

Consider first the fully observed case, i.e., every $a_{i,j} = 0$ or 1. Suppose also that the graph is already ideally clustered—i.e., there is a partition of the nodes such that there is no edge between clusters, and each cluster is a clique. In this case, the matrix $\mathbf{A} + \mathbf{I}$ is now a *low-rank* matrix, with the rank equal to the number of clusters. This can be seen by noticing that if we re-order the rows and columns so that clusters appear together, the result would be a *block-diagonal* matrix, with each block being an all-ones submatrix—and thus rank one. Of course, this re-ordering does not change the rank of the matrix, and hence $\mathbf{A} + \mathbf{I}$ is exactly low-rank.

Consider now any given graph, still fully observed. In light of the above, we are looking for a decomposition of its $\mathbf{A} + \mathbf{I}$ into a low-rank part \mathbf{K}^* (of block-diagonal all-ones, one block for each cluster) and a remaining \mathbf{B}^* (the disagreements), such that the number of non-zero entries in \mathbf{B}^* is as small as possible; i.e., \mathbf{B}^* is sparse. Finally, the problem we look at is recovery of the best \mathbf{K}^* when we do not observe all entries of $\mathbf{A} + \mathbf{I}$. The idea is depicted in Figure 1.

1. In practice, one might be able to use the “failed” output with rounding as an approximate solution. In this paper, we focus on the performance of the unrounded algorithm.




Figure 1: The adjacency matrix of a graph drawn from the planted partition model before and after proper reordering (i.e., clustering) of the nodes. The figure on the right is indicative of the matrix as a superposition of a sparse matrix and a low-rank block diagonal one.

We propose to perform the matrix splitting using convex optimization (Chandrasekaran et al., 2011; Candès et al., 2011). Our approach consists of dropping any additional structural requirements, and just looking for a decomposition of the given $\mathbf{A} + \mathbf{I}$ as the sum of a sparse matrix \mathbf{B} and a low-rank matrix \mathbf{K} . Recall that Ω_{obs} is the set of observed entries, i.e., the set of elements of \mathbf{A} that are known to be 0 or 1; we use the following convex program:

$$\begin{aligned} \min_{\mathbf{B}, \mathbf{K}} \quad & \lambda \|\mathbf{B}\|_1 + \|\mathbf{K}\|_* \\ \text{s.t.} \quad & \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{B} + \mathbf{K}) = \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{A} + \mathbf{I}). \end{aligned} \tag{1}$$

Here, for any matrix \mathbf{M} , the term $\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{M})$ keeps all elements of \mathbf{M} in Ω_{obs} unchanged, and sets all other elements to 0; the constraints thus state that the sparse and low-rank matrix should in sum be consistent with the observed entries of $\mathbf{A} + \mathbf{I}$. The term $\|\mathbf{B}\|_1 = \sum_{i,j} |b_{i,j}|$ is the ℓ_1 norm of the entries of the matrix \mathbf{B} , which is well-known to be a convex surrogate for the number of non-zero entries $\|\mathbf{B}\|_0$. The second term $\|\mathbf{K}\|_* = \sum_s \sigma_s(\mathbf{K})$ is the nuclear norm (also known as the trace norm), defined as the sum of the singular values of \mathbf{K} . This has been shown to be the tightest convex surrogate for the rank function for matrices with unit spectral norm (Fazel, 2002). Thus our objective function is a convex surrogate for the (natural) combinatorial objective $\lambda\|\mathbf{B}\|_0 + \text{rank}(\mathbf{K})$. The optimization problem (1) is, in fact, a semidefinite program (SDP) (Chandrasekaran et al., 2011).

We remark on the above formulation. (a) This formulation does not require specifying the number of clusters; this parameter is effectively learned from the data. The tradeoff parameter λ is artificial and can be easily determined: since any desired \mathbf{K}^* has trace exactly equal to n , we simply choose the smallest λ such that the trace of the optimal solution is at least n . This can be done by, e.g., bisection, which is described below. (b) It is possible to obtain tighter convex relaxations by adding more constraints, such as the diagonal entry constraints $k_{i,i} = 1, \forall i$, the positive semi-definite constraint $\mathbf{K} \succeq 0$, or even the triangular inequalities $k_{i,j} + k_{j,k} - k_{i,k} \leq 1$. Indeed, this is done by Mathieu and Schudy (2010). Note that the guarantees for our formulation (to be presented in the next subsection) automatically imply guarantees for any other tighter relaxations. We choose to focus on our (looser) formulation for two reasons. First, and most importantly, even with the extra

constraints, Mathieu and Schudy (2010) do not deliver better exact recovery guarantees (cf. Table 1). In fact, we show in Section 2.3 that our results are near optimal in some important regimes, so tighter relaxations do not seem to provide additional benefits in exact recovery. Second, our formulation can be solved efficiently using existing Augmented Lagrangian Multiplier methods (Lin et al., 2009). This is no longer the case with the $\Theta(n^3)$ triangle inequality constraints enforced by Mathieu and Schudy (2010), and solving it as a standard SDP is only feasible for small graphs.

We are interested in the case when the convex program (1) produces an optimal solution \mathbf{K} that is a block-diagonal matrix and corresponds to an ideal clustering.

Definition 1 (Validity) *The convex program (1) is said to produce a valid output if the low-rank matrix part \mathbf{K} of the optimal solution corresponds to a graph of disjoint cliques; i.e., its rows and columns can be re-ordered to yield a block-diagonal matrix with all ones for each block.*

Validity of a given \mathbf{K} can be easily checked via elementary re-ordering operations.² Our first simple but useful insight is that whenever the convex program (1) yields a valid solution, it is the disagreement minimizer.

Theorem 2 *For any $\lambda > 0$, if the solution of (1) is valid, then it is the clustering that minimizes the number of observed disagreements.*

Our complete clustering procedure is given as Algorithm 1. It takes the adjacency matrix of the graph \mathbf{A} and outputs either the optimal clustering or declares failure. Setting the parameter λ is done via binary search. The initial value of λ is not crucial; we use $\lambda = \frac{1}{32\sqrt{\bar{p}_0 n}}$ based on our theoretical analysis in the next sub-section, where \bar{p}_0 is the empirical fraction of observed pairs. To solve the optimization problem (1), we use the fast algorithm developed by Lin et al. (2009), which is tailored for matrix splitting and takes advantage of the sparsity of the observations. By Theorem 2, whenever the algorithm results in a valid \mathbf{K} , we have found the optimal clustering.

Algorithm 1 Optimal-Cluster(\mathbf{A})

```

 $\lambda \leftarrow \frac{1}{32\sqrt{\bar{p}_0 n}}$ 
while not terminated do
  Solve (1) to obtain the solution ( $\mathbf{B}, \mathbf{K}$ )
  if  $\mathbf{K}$  is valid then
    Output the clustering in  $\mathbf{K}$  and EXIT.
  else if  $\text{trace}(\mathbf{K}) > n$  then
     $\lambda \leftarrow \lambda/2$ 
  else if  $\text{trace}(\mathbf{K}) < n$  then
     $\lambda \leftarrow 2\lambda$ 
  end if
end while
Declare Failure.

```

2. If we re-order a valid \mathbf{K} such that identical rows and columns appear together, it will become block-diagonal.

2.2 Performance Analysis

For the main analytical contribution of this paper, we provide conditions under which the above algorithm will find the clustering that minimizes the number of disagreements among the observed entries. In particular, we characterize its performance under the standard and classical planted partition/stochastic block model with partial observations, which we now describe.

Definition 3 (Planted Partition Model with Partial Observations) *Suppose that n nodes are partitioned into r clusters, each of size at least K_{\min} . Let \mathbf{K}^* be the low-rank matrix corresponding to this clustering (as described above). The adjacency matrix \mathbf{A} of the graph is generated as follows: for each pair of nodes (i, j) in the same cluster, $a_{i,j} = ?$ with probability $1 - p_0$, $a_{i,j} = 1$ with probability p_0p , or $a_{i,j} = 0$ otherwise, independent of all others; similarly, for (i, j) in different clusters, $a_{i,j} = ?$ with probability $1 - p_0$, $a_{i,j} = 1$ with probability p_0q , or $a_{i,j} = 0$ otherwise.*

Under the above model, the graph is observed at locations chosen at random with probability p_0 . In expectation a fraction of $1 - p$ of the in-cluster observations are disagreements; similarly, the fraction of disagreements in the across-cluster observations is q . Let $\mathbf{B}^* = \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{A} + \mathbf{I} - \mathbf{K}^*)$ be the matrix of observed disagreements for the original clustering; note that the support of \mathbf{B}^* is contained in Ω_{obs} . The following theorem provides a sufficient condition for our algorithm to recover the original clustering $(\mathbf{B}^*, \mathbf{K}^*)$ with high probability. Combined with Theorem 2, it also shows that under the same condition, the original clustering is disagreement minimizing with high probability.

Theorem 4 *Let $\tau = \max\{1 - p, q\}$. There exist universal positive constants c and C such that, with probability at least $1 - cn^{-10}$, the original clustering $(\mathbf{B}^*, \mathbf{K}^*)$ is the unique optimal solution of (1) with $\lambda = \frac{1}{32\sqrt{np_0}}$ provided that*

$$p_0(1 - 2\tau)^2 \geq C \frac{n \log^2 n}{K_{\min}^2}. \quad (2)$$

Note that the quantity τ is (an upper bound of) the probability of having a disagreement, and $1 - 2\tau$ is (a lower bound of) the density gap $p - q$. The sufficient condition in the theorem is given in terms of the three parameters that define problem: the minimum cluster size K_{\min} , the density gap $1 - 2\tau$, and the observation probability p_0 . We remark on these parameters.

- *Minimum cluster size K_{\min} .* Since the left hand side of the condition (2) in Theorem 4 is no more than 1, it imposes a lower-bound $K_{\min} = \tilde{\Omega}(\sqrt{n})$ on the cluster sizes. This means that our method can handle a growing number $\tilde{O}(\sqrt{n})$ of clusters. The lower-bound on K_{\min} is attained when $1 - 2\tau$ and p_0 are both $\Theta(1)$, i.e., not decreasing as n grows. Note that all relevant works require a lower-bound at least as strong as ours (cf. Table 1).
- *Density gap $1 - 2\tau$.* When $p_0 = \Theta(1)$, our result allows this gap to be vanishingly small, i.e., $\tilde{\Omega}\left(\frac{\sqrt{n}}{K_{\min}}\right)$, where a larger K_{\min} allows for a smaller gap. As we mentioned in Section 1.2, this matches the best available results (cf. Table 1), including those in Mathieu and Schudy (2010) and Oymak and Hassibi (2011), which use tighter convex relaxations that are more computationally demanding. We note that directly applying existing results in the low-rank-plus-sparse literature (Candès et al., 2011; Li, 2013) leads to weaker results, where the gap is bounded below by a constant.

- *Observation probability* p_0 . When $1 - 2\tau = \Theta(1)$, our result only requires a vanishing fraction of observations, i.e., p_0 can be as small as $\tilde{\Theta}\left(\frac{n}{K_{\min}^2}\right)$; a larger K_{\min} allows for a smaller p_0 . As mentioned in Section 1.2, this scaling is better than prior results we know of.
- *Tradeoffs*. A novel aspect of our result is that it shows an explicit tradeoff between the observation probability p_0 and the density gap $1 - 2\tau$. The left hand side of (2) is linear in p_0 and quadratic in $1 - 2\tau$. This means if the number of observations is four times larger, then we can handle a 50% smaller density gap. Moreover, p_0 can go to zero quadratically faster than $1 - 2\tau$. Consequently, treating missing observations as disagreements would lead to quadratically weaker results. This agrees with the intuition that handling missing entries with known locations is easier than correcting disagreements whose locations are unknown.

We would like to point out that our algorithm has the capability to handle outliers. Suppose there are some isolated nodes which do not belong to any cluster, and they connect to each other and each node in the clusters with probability at most τ , with τ obeying the condition (2) in Theorem 4. Our algorithm will classify all these edges as disagreements, and hence automatically reveal the identity of each outlier. In the output of our algorithm, the low rank part \mathbf{K} will have all zeros in the columns and rows corresponding to outliers—all their edges will appear in the disagreement matrix \mathbf{B} .

2.3 Lower Bounds

We now discuss the tightness of Theorem 4. Consider first the case where $K_{\min} = \Theta(n)$, which means there are a constant number of clusters. We establish a fundamental lower bound on the density gap $1 - 2\tau$ and the observation probability p_0 that are required for *any* algorithm to correctly recover the clusters.

Theorem 5 *Under the planted partition model with partial observations, suppose the true clustering is chosen uniformly at random from all possible clusterings with equal cluster size K . If $K = \Theta(n)$ and $\tau = 1 - p = q > 1/100$, then for any algorithm to correctly identify the clusters with probability at least $\frac{3}{4}$, we need*

$$p_0(1 - 2\tau)^2 \geq C \frac{1}{n},$$

where $C > 0$ is an absolute constant.

Theorem 5 generalizes a similar result in Chaudhuri et al. (2012), which does not consider partial observations. The theorem applies to any algorithm regardless of its computational complexity, and characterizes the fundamental tradeoff between p_0 and $1 - 2\tau$. It shows that when $K_{\min} = \Theta(n)$, the requirement for $1 - 2\tau$ and p_0 in Theorem 4 is optimal up to logarithmic factors, and cannot be significantly improved by using more complicated methods.

For the general case with $K_{\min} = O(n)$, only part of the picture is known. Using non-rigorous arguments, Decelle et al. (2011) show that $1 - 2\tau \gtrsim \frac{\sqrt{n}}{K_{\min}}$ is necessary when $\tau = \Theta(1)$ and the graph is fully observed; otherwise recovery is impossible or computationally hard. According to this lower-bound, our requirement on the density gap $1 - 2\tau$ is probably tight (up to log factors) for all K_{\min} . However, a rigorous proof of this claim is still lacking, and seems to be a difficult problem. Similarly, the tightness of our condition on p_0 and the tradeoff between p_0 and τ is also unclear in this regime.

3. Proofs

In this section, we prove Theorems 2 and 4. The proof of Theorem 5 is deferred to Appendix B.

3.1 Proof of Theorem 2

We first prove Theorem 2, which says that if the optimization problem (1) produces a valid matrix, i.e., one that corresponds to a clustering of the nodes, then this is the disagreement minimizing clustering. Consider the following non-convex optimization problem

$$\begin{aligned} \min_{\mathbf{B}, \mathbf{K}} \quad & \lambda \|\mathbf{B}\|_1 + \|\mathbf{K}\|_* \\ \text{s.t.} \quad & \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{B} + \mathbf{K}) = \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{I} + \mathbf{A}), \\ & \mathbf{K} \text{ is valid,} \end{aligned} \tag{3}$$

and let (\mathbf{B}, \mathbf{K}) be any feasible solution. Since \mathbf{K} represents a valid clustering, it is positive semidefinite and has all ones along its diagonal. Therefore, it obeys $\|\mathbf{K}\|_* = \text{trace}(\mathbf{K}) = n$. On the other hand, because both $\mathbf{K} - \mathbf{I}$ and \mathbf{A} are adjacency matrices, the entries of $\mathbf{B} = \mathbf{I} + \mathbf{A} - \mathbf{K}$ in Ω_{obs} must be equal to -1 , 1 or 0 (i.e., it is a disagreement matrix). Clearly any optimal \mathbf{B} must have zeros at the entries in Ω_{obs}^c . Hence $\|\mathbf{B}\|_1 = \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{B})\|_0$ when \mathbf{K} is valid. We thus conclude that the above optimization problem (3) is equivalent to minimizing $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{B})\|_0$ subject to the same constraints. This is exactly the minimization of the number of disagreements on the observed edges. Now notice that (1) is a relaxed version of (3). Therefore, if the optimal solution of (1) is valid and thus feasible to (3), then it is also optimal to (3), the disagreement minimization problem.

3.2 Proof of Theorem 4

We now turn to the proof of Theorem 4, which provides guarantees for when the convex program (1) recovers the true clustering $(\mathbf{B}^*, \mathbf{K}^*)$.

3.2.1 PROOF OUTLINE AND PRELIMINARIES

We overview the main steps in the proof of Theorem 4; details are provided in Sections 3.2.2–3.2.4 to follow. We would like to show that the pair $(\mathbf{B}^*, \mathbf{K}^*)$ corresponding to the true clustering is the unique optimal solution to our convex program (1). This involves the following three steps.

Step 1: We show that it suffices to consider an equivalent model for the observation and disagreements. This model is easier to handle, especially when the observation probability and density gap are vanishingly small, which is the regime of interest in this paper.

Step 2: We write down the sub-gradient based first-order sufficient conditions that need to be satisfied for $(\mathbf{B}^*, \mathbf{K}^*)$ to be the unique optimum of (1). In our case, this involves showing the existence of a matrix \mathbf{W} —the *dual certificate*—that satisfies certain properties. This step is technical—requiring us to deal with the intricacies of sub-gradients since our convex function is not smooth—but otherwise standard. Luckily for us, this has been done previously (Chandrasekaran et al., 2011; Candès et al., 2011; Li, 2013).

Step 3: Using the assumptions made on the true clustering \mathbf{K}^* and its disagreements \mathbf{B}^* , we construct a candidate dual certificate \mathbf{W} that meets the requirements in step 2, and thus certify $(\mathbf{B}^*, \mathbf{K}^*)$ as being the unique optimum.

The crucial Step 3 is where we go beyond the existing literature on matrix splitting (Chandrasekaran et al., 2011; Candès et al., 2011; Li, 2013). These results assume the observation probability and/or density gap is at least a constant, and hence do not apply to our setting. Here we provide a refined analysis, which leads to better performance guarantees than those that could be obtained via a direct application of existing sparse and low-rank matrix splitting results.

Next, we introduce some notations used in the rest of the proof of the theorem. The following definitions related to \mathbf{K}^* are standard. By symmetry, the SVD of \mathbf{K}^* has the form $\mathbf{U}\Sigma\mathbf{U}^T$, where $\mathbf{U} \in \mathbb{R}^{n \times r}$ contains the singular vectors of \mathbf{K}^* . We define the subspace

$$\mathcal{T} \triangleq \{\mathbf{U}\mathbf{X}^T + \mathbf{Y}\mathbf{U}^T : \mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times r}\},$$

which is spanned of all matrices that share either the same column space or the same row space as \mathbf{K}^* . For any matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$, its orthogonal projection to the space \mathcal{T} is given by $\mathcal{P}_{\mathcal{T}}(\mathbf{M}) = \mathbf{U}\mathbf{U}^T\mathbf{M} + \mathbf{M}\mathbf{U}\mathbf{U}^T - \mathbf{U}\mathbf{U}^T\mathbf{M}\mathbf{U}\mathbf{U}^T$. The projection onto \mathcal{T}^\perp , the complement orthogonal space of \mathcal{T} , is given by $\mathcal{P}_{\mathcal{T}^\perp}(\mathbf{M}) = \mathbf{M} - \mathcal{P}_{\mathcal{T}}(\mathbf{M})$.

The following definitions are related to \mathbf{B}^* and partial observations. Let $\Omega^* = \{(i, j) : b_{i,j}^* \neq 0\}$ be the set of matrix entries corresponding to the disagreements. Recall that Ω_{obs} is the set of observed entries. For any matrix \mathbf{M} and entry set Ω_0 , we let $\mathcal{P}_{\Omega_0}(\mathbf{M}) \in \mathbb{R}^{n \times n}$ be the matrix obtained from \mathbf{M} by setting all entries not in the set Ω_0 to zero. We write $\Omega_0 \sim \text{Ber}_0(p)$ if the entry set Ω_0 does not contain the diagonal entries, and each pair (i, j) and (j, i) ($i \neq j$) is contained in Ω_0 with probability p , independent all others; $\Omega_0 \sim \text{Ber}_1(p)$ is defined similarly except that Ω_0 contains all the diagonal entries. Under our partially observed planted partition model, we have $\Omega_{\text{obs}} \sim \text{Ber}_1(p_0)$ and $\Omega^* \sim \text{Ber}_0(\tau)$.

Several matrix norms are used. $\|\mathbf{M}\|$ and $\|\mathbf{M}\|_F$ represent the spectral and Frobenius norms of the matrix \mathbf{M} , respectively, and $\|\mathbf{M}\|_\infty \triangleq \max_{i,j} |m_{i,j}|$ is the matrix infinity norm.

3.2.2 STEP 1: EQUIVALENT MODEL FOR OBSERVATIONS AND DISAGREEMENTS

It is easy show that increasing p or decreasing q can only make the probability of success higher, so without loss of generality we assume $1 - p = q = \tau$. Observe that the probability of success is completely determined by the distribution of $(\Omega_{\text{obs}}, \mathbf{B}^*)$ under the planted partition model with partial observations. The first step is to show that it suffices to consider an equivalent model for generating $(\Omega_{\text{obs}}, \mathbf{B}^*)$, which results in the same distribution but is easier to handle. This is in the same spirit as Candès et al. (2011, Theorems 2.2 and 2.3) and Li (2013, Section 4.1). In particular, we consider the following procedure:

1. Let $\Gamma \sim \text{Ber}_1(p_0(1 - 2\tau))$, and $\Omega \sim \text{Ber}_0\left(\frac{2p_0\tau}{1 - p_0 + 2p_0\tau}\right)$. Let $\Omega_{\text{obs}} = \Gamma \cup \Omega$.
2. Let \mathbf{S} be a symmetric random matrix whose upper-triangular entries are independent and satisfy $\mathbb{P}(s_{i,j} = 1) = \mathbb{P}(s_{i,j} = -1) = \frac{1}{2}$.
3. Define $\Omega' \subseteq \Omega$ as $\Omega' = \{(i, j) : (i, j) \in \Omega, s_{i,j} = 1 - 2k_{i,j}^*\}$. In other words, Ω' is the entries of \mathbf{S} whose signs are consistent with a disagreement matrix.
4. Define $\Omega^* = \Omega' \setminus \Gamma$, and $\tilde{\Gamma} = \Omega_{\text{obs}} \setminus \Omega^*$.
5. Let $\mathbf{B}^* = \mathcal{P}_{\Omega^*}(\mathbf{S})$.

It is easy to verify that $(\Omega_{\text{obs}}, \mathbf{B}^*)$ has the same distribution as in the original model. In particular, we have $\mathbb{P}[(i, j) \in \Omega_{\text{obs}}] = p_0$, $\mathbb{P}[(i, j) \in \Omega^*, (i, j) \in \Omega_{\text{obs}}] = p_0\tau$ and $\mathbb{P}[(i, j) \in \Omega^*, (i, j) \notin \Omega_{\text{obs}}] = 0$, and observe that given $\mathbf{K}^*, \mathbf{B}^*$ is completely determined by its support Ω^* .

The advantage of the above model is that Γ and Ω are independent of each other, and \mathbf{S} has random signed entries. This facilitates the construction of the dual certificate, especially in the regime of vanishing p_0 and $(\frac{1}{2} - \tau)$ considered in this paper. We use this equivalent model in the rest of the proof.

3.2.3 STEP 2: SUFFICIENT CONDITIONS FOR OPTIMALITY

We state the first-order conditions that guarantee $(\mathbf{B}^*, \mathbf{K}^*)$ to be the unique optimum of (1) with high probability. Here and henceforth, by *with high probability* we mean with probability at least $1 - cn^{-10}$ for some universal constant $c > 0$. The following lemma follows from Theorem 4.4 in Li (2013) and the discussion thereafter.

Lemma 6 (Optimality Condition) *Suppose $\left\| \frac{1}{(1-2\tau)p_0} \mathcal{P}_{\mathcal{T}} \mathcal{P}_{\Gamma} \mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}} \right\| \leq \frac{1}{2}$. Then $(\mathbf{B}^*, \mathbf{K}^*)$ is the unique optimal solution to (1) with high probability if there exists $\mathbf{W} \in \mathbb{R}^{n \times n}$ such that*

1. $\left\| \mathcal{P}_{\mathcal{T}}(\mathbf{W} + \lambda \mathcal{P}_{\Omega} \mathbf{S} - \mathbf{U} \mathbf{U}^{\top}) \right\|_F \leq \frac{\lambda}{n^2}$,
2. $\left\| \mathcal{P}_{\mathcal{T}^{\perp}}(\mathbf{W} + \lambda \mathcal{P}_{\Omega} \mathbf{S}) \right\| \leq \frac{1}{4}$,
3. $\mathcal{P}_{\Gamma^c}(\mathbf{W}) = 0$,
4. $\left\| \mathcal{P}_{\Gamma}(\mathbf{W}) \right\|_{\infty} \leq \frac{\lambda}{4}$.

Lemma 9 in the appendix guarantees that the condition $\left\| \frac{1}{(1-2\tau)p_0} \mathcal{P}_{\mathcal{T}} \mathcal{P}_{\Gamma} \mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}} \right\| \leq \frac{1}{2}$ is satisfied with high probability under the assumption of Theorem 4. It remains to show the existence of a desired dual certificate \mathbf{W} which satisfies the four conditions in Lemma 6 with high probability.

3.2.4 STEP 3: DUAL CERTIFICATE CONSTRUCTION

We use a variant of the so-called *golfing scheme* (Candès et al., 2011; Gross, 2011) to construct \mathbf{W} . Our application of golfing scheme, as well as its analysis, is different from previous work and leads to stronger guarantees. In particular, we go beyond existing results by allowing the fraction of observed entries and the density gap to be vanishing.

By definition in Section 3.2.2, Γ obeys $\Gamma \sim \text{Ber}_1(p_0(1-2\tau))$. Observe that Γ may be considered to be generated by $\Gamma = \bigcup_{1 \leq k \leq k_0} \Gamma_k$, where the sets $\Gamma_k \sim \text{Ber}_1(t)$ are independent; here the parameter t obeys $p_0(1-2\tau) = 1 - (1-t)^{k_0}$, and k_0 is chosen to be $\lceil 5 \log n \rceil$. This implies $t \geq p_0(1-2\tau)/k_0 \geq C_0 \frac{n \log n}{K_{\min}^2}$ for some constant C_0 , with the last inequality holds under the assumption of Theorem 4. For any random entry set $\Omega_0 \sim \text{Ber}_1(\rho)$, define the operator $\mathcal{R}_{\Omega_0} : \mathbb{R}^{n \times n} \mapsto \mathbb{R}^{n \times n}$ by

$$\mathcal{R}_{\Omega_0}(\mathbf{M}) = \sum_{i=1}^n m_{i,i} e_i e_i^T + \rho^{-1} \sum_{1 \leq i < j \leq n} \delta_{ij} m_{i,j} (e_i e_j^T + e_j e_i^T),$$

where δ_{ij} is the indicator random variable with $\delta_{ij} = 1$ if $(i, j) \in \Omega_0$ and 0 otherwise, and e_i is the i -th standard basis in $\mathbb{R}^{n \times n}$, i.e., the column vector with 1 in its i -th entry and 0 elsewhere.

We now define our dual certificate. Let $\mathbf{W} = \mathbf{W}_{k_0}$, where \mathbf{W}_{k_0} is defined recursively by setting $\mathbf{W}_0 = 0$ and for all $k = 1, 2, \dots, k_0$,

$$\mathbf{W}_k = \mathbf{W}_{k-1} + \mathcal{R}_{\Gamma_k} \mathcal{P}_{\mathcal{T}} (\mathbf{U}\mathbf{U}^T - \lambda \mathcal{P}_{\mathcal{T}}(\mathcal{P}_{\Omega}(\mathbf{S})) - \mathbf{W}_{k-1}).$$

Clearly the equality condition in Lemma 6 is satisfied. It remains to show that \mathbf{W} also satisfies the inequality conditions with high probability. The proof makes use of the technical Lemmas 9–12 given in the appendix. For convenience of notation, we define the quantity $\Delta_k = \mathbf{U}\mathbf{U}^T - \lambda \mathcal{P}_{\mathcal{T}}(\mathcal{P}_{\Omega}(\mathbf{S})) - \mathcal{P}_{\mathcal{T}}(\mathbf{W}_k)$, and use the notation

$$\prod_{i=1}^k (\mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}} \mathcal{R}_{\Gamma_i} \mathcal{P}_{\mathcal{T}}) = (\mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}} \mathcal{R}_{\Gamma_k} \mathcal{P}_{\mathcal{T}}) \cdots (\mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}} \mathcal{R}_{\Gamma_1} \mathcal{P}_{\mathcal{T}}),$$

where the order of multiplication is important. Observe that by construction of \mathbf{W} , we have

$$\Delta_k = \prod_{i=1}^k (\mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}} \mathcal{R}_{\Gamma_i} \mathcal{P}_{\mathcal{T}}) (\mathbf{U}\mathbf{U}^T - \lambda \mathcal{P}_{\mathcal{T}} \mathcal{P}_{\Omega}(\mathbf{S})), k = 1, \dots, k_0, \quad (4)$$

$$\mathbf{W}_{k_0} = \sum_{k=1}^{k_0} \mathcal{R}_{\Gamma_k} \Delta_{k-1}. \quad (5)$$

We are ready to prove that \mathbf{W} satisfies inequalities 1, 2 and 4 in Lemma 6.

Inequality 1: Thanks to (4), we have the following geometric convergence :

$$\begin{aligned} \left\| \mathcal{P}_{\mathcal{T}}(\mathbf{W} + \lambda \mathcal{P}_{\Omega} \mathbf{S} - \mathbf{U}\mathbf{U}^T) \right\|_F &= \|\Delta_{k_0}\|_F \\ &\leq \left(\prod_{k=1}^{k_0} \|\mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}} \mathcal{R}_{\Gamma_k} \mathcal{P}_{\mathcal{T}}\| \right) \|\mathbf{U}\mathbf{U}^T - \lambda \mathcal{P}_{\mathcal{T}} \mathcal{P}_{\Omega}(\mathbf{S})\|_F \\ &\stackrel{(a)}{\leq} e^{-k_0} (\|\mathbf{U}\mathbf{U}^T\|_F + \lambda \|\mathcal{P}_{\mathcal{T}} \mathcal{P}_{\Omega}(\mathbf{S})\|_F) \\ &\stackrel{(b)}{\leq} n^{-5} (n + \lambda \cdot n) \leq (1 + \lambda) n^{-4} \stackrel{(c)}{\leq} \frac{1}{2n^2} \lambda. \end{aligned}$$

Here, the inequality (a) follows Lemma 9 with $\epsilon_1 = e^{-1}$, (b) follows from our choices of λ and k_0 and the fact that $\|\mathcal{P}_{\mathcal{T}} \mathcal{P}_{\Omega}(\mathbf{S})\|_F \leq \|\mathcal{P}_{\Omega}(\mathbf{S})\|_F \leq n$, and (c) holds under the assumption $\lambda \geq \frac{1}{32\sqrt{n}}$ in the theorem.

Inequality 4: We have

$$\|\mathcal{P}_{\Gamma}(\mathbf{W})\|_{\infty} = \|\mathcal{P}_{\Gamma}(\mathbf{W}_{k_0})\|_{\infty} \leq \sum_{k=1}^{k_0} \|\mathcal{R}_{\Gamma_k} \Delta_{k-1}\|_{\infty} \leq t^{-1} \sum_{k=1}^{k_0} \|\Delta_{k-1}\|_{\infty},$$

where the first inequality follows from (5) and the triangle inequality. We proceed to obtain

$$\begin{aligned}
 \sum_{k=1}^{k_0} \|\Delta_{k-1}\|_\infty &\stackrel{(a)}{=} \sum_{k=1}^{k_0} \left\| \prod_{i=1}^{k-1} (\mathcal{P}_\mathcal{T} - \mathcal{P}_\mathcal{T} \mathcal{R}_{\Gamma_i} \mathcal{P}_\mathcal{T}) (\mathbf{U} \mathbf{U}^T - \lambda \mathcal{P}_\mathcal{T} \mathcal{P}_\Omega(\mathbf{S})) \right\|_\infty \\
 &\stackrel{(b)}{\leq} \sum_{k=1}^{k_0} \left(\frac{1}{2}\right)^k \|\mathbf{U} \mathbf{U}^T - \lambda \mathcal{P}_\mathcal{T} \mathcal{P}_\Omega(\mathbf{S})\|_\infty \\
 &\stackrel{(c)}{\leq} \frac{1}{K_{\min}} + \lambda \sqrt{\frac{p_0 n \log n}{K_{\min}^2}}, \tag{6}
 \end{aligned}$$

where (a) follows from (4), (b) follows from Lemma 11 and (c) follows from Lemma 12. It follows that

$$\begin{aligned}
 \|\mathcal{P}_\Gamma(\mathbf{W})\|_\infty &\leq \frac{1}{t} \left(\frac{1}{K_{\min}} + \frac{n \log n}{K_{\min}^2} \lambda \right) \\
 &\leq \frac{k_0}{p_0(1-2\tau)} \left(\frac{1}{K_{\min}} + \lambda \sqrt{\frac{p_0 n \log n}{K_{\min}^2}} \right) \leq \frac{1}{4} \lambda,
 \end{aligned}$$

where the last inequality holds under the assumption of Theorem 4

Inequality 2: Observe that by the triangle inequality, we have

$$\|\mathcal{P}_{\mathcal{T}^\perp}(\mathbf{W} + \lambda \mathcal{P}_\Omega(\mathbf{S}))\| \leq \lambda \|\mathcal{P}_{\mathcal{T}^\perp}(\mathcal{P}_\Omega(\mathbf{S}))\| + \|\mathcal{P}_{\mathcal{T}^\perp}(\mathbf{W}_{k_0})\|.$$

For the first term, standard results on the norm of a matrix with i.i.d. entries (e.g., see Vershynin 2010) give

$$\lambda \|\mathcal{P}_{\mathcal{T}^\perp}(\mathcal{P}_\Omega(\mathbf{S}))\| \leq \lambda \|\mathcal{P}_\Omega(\mathbf{S})\| \leq \frac{1}{32\sqrt{p_0 n}} \cdot 4 \sqrt{\frac{2p_0 \tau n}{1-p_0+2p_0\tau}} \leq \frac{1}{8}$$

It remains to show that the second term is bounded by $\frac{1}{8}$. To this end, we observe that

$$\begin{aligned}
 \|\mathcal{P}_{\mathcal{T}^\perp}(\mathbf{W}_{k_0})\| &\stackrel{(a)}{=} \sum_{k=1}^{k_0} \|\mathcal{P}_{\mathcal{T}^\perp}(\mathcal{R}_{\Gamma_k} \Delta_{k-1} - \Delta_{k-1})\| \\
 &\leq \sum_{k=1}^{k_0} \|(\mathcal{R}_{\Gamma_k} - \mathcal{I}) \Delta_{k-1}\| \\
 &\stackrel{(b)}{\leq} C \sqrt{\frac{n \log n}{t}} \sum_{k=1}^{k_0} \|\Delta_{k-1}\|_\infty \\
 &\stackrel{(c)}{\leq} C \sqrt{\frac{k_0 n \log n}{p_0(1-2\tau)}} \left(\frac{1}{K_{\min}} + \lambda \sqrt{\frac{p_0 n \log n}{K_{\min}^2}} \right) \leq \frac{1}{8},
 \end{aligned}$$

where in (a) we use (5) and the fact that $\Delta_k \in \mathcal{T}$, (b) follows from Lemma 10, and (c) follows from (6). This completes the proof of Theorem 4.

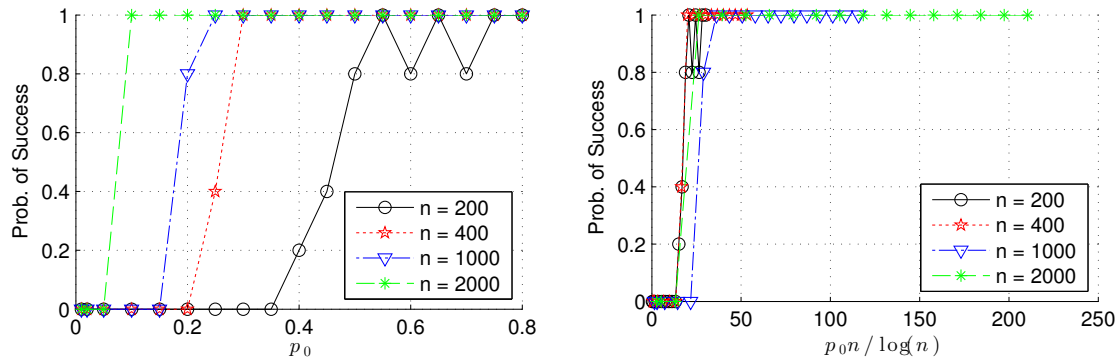


Figure 2: Simulation results verifying the performance of our algorithm as a function of the observation probability p_0 and the graph size n . The left pane shows the probability of successful recovery under different p_0 and n with fixed $\tau = 0.2$ and $K_{\min} = n/4$; each point is an average over 5 trials. After proper rescaling of the x-axis, the curves align as shown in the right pane, indicating a good match with our theoretical results.

4. Experimental Results

We explore via simulation the performance of our algorithm as a function of the values of the model parameters (n, K_{\min}, p_0, τ) . We see that the performance matches well with the theory.

In the experiment, each test case is constructed by generating a graph with n nodes divided into clusters of equal size K_{\min} , and then placing a disagreement on each pair of node with probability τ independently. Each node pair is then observed with probability p_0 . We then run Algorithm 1, where the optimization problem (1) is solved using the fast algorithm in Lin et al. (2009). We check if the algorithm successfully outputs a solution that equals to the underlying true clusters. In the first set of experiments, we fix $\tau = 0.2$ and $K_{\min} = n/4$ and vary p_0 and n . For each (p_0, n) , we repeat the experiment for 5 times and plot the probability of success in the left pane of Figure 2.

One observes that our algorithm has better performance with larger p_0 and n , and the success probability exhibits a phase transition. Theorem 4 predicts that, with τ fixed and $K_{\min} = n/4$, the transition occurs at $p_0 \propto \frac{n \log^2 n}{K_{\min}^2} \propto \frac{\log^2 n}{n}$; in particular, if we plot the success probability against the control parameter $\frac{p_0 n}{\log^2 n}$, all curves should align with each other. Indeed, this is precisely what we see in the right pane of Figure 2 where we use $\frac{p_0 n}{\log n}$ as the control parameter. This shows that Theorem 4 gives the correct scaling between p_0 and n up to an extra log factor.

In a similar fashion, we run another three sets of experiments with the following settings: (1) $n = 1000$ and $\tau = 0.2$ with varying (p_0, K_{\min}) ; (2) $K_{\min} = n/4$ and $p_0 = 0.2$ with varying (τ, n) ; (3) $n = 1000$ and $p_0 = 0.6$ with varying (τ, K_{\min}) . The results are shown in Figures 3, 4 and 5; note that each x -axis corresponds to a control parameter chosen according to the scaling predicted by Theorem 4. Again we observe that all the curves roughly align, indicating a good match with the theory. In particular, by comparing Figures 2 and 4 (or Figures 3 and 5), one verifies the quadratic tradeoff between observations and disagreements (i.e., p_0 vs. $1 - 2\tau$) as predicted by Theorem 4.

Finally, we compare the performance of our method with spectral clustering, a popular method for graph clustering. For spectral clustering, we first impute the missing entries of the adjacency




Figure 3: Simulation results verifying the performance of our algorithm as a function of the observation probability p_0 and the cluster size K_{\min} , with $n = 1000$ and $\tau = 0.2$ fixed.

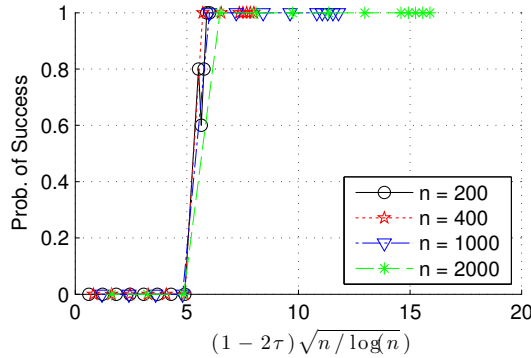


Figure 4: Simulation results verifying the performance of our algorithm as a function of the disagreement probability τ and the graph size n , with $p_0 = 0.2$ and $K_{\min} = n/4$ fixed.

matrix with either zeros or random $1/0$'s. We then compute the first k principal components of the adjacency matrix, and run k -means clustering on the principal components (von Luxburg, 2007); here we set k equal to the number of clusters. The adjacency matrix is generated in the same fashion as before using the parameters $n = 2000$, $K_{\min} = 200$ and $\tau = 0.1$. We vary the observation probability p_0 and plot the success probability in Figure 6. It can be observed that our method outperforms spectral clustering with both imputation schemes; in particular, it requires fewer observations.

5. Conclusion

We proposed a convex optimization formulation, based on a reduction to decomposing low-rank and sparse matrices, to address the problem of clustering partially observed graphs. We showed that under a wide range of parameters of the planted partition model with partial observations, our method is guaranteed to find the optimal (disagreement-minimizing) clustering. In particular, our method succeeds under higher levels of noise and/or missing observations than existing methods in




Figure 5: Simulation results verifying the performance of our algorithm as a function of the disagreement probability τ and the cluster size K_{\min} , with $n = 1000$ and $p_0 = 0.6$ fixed.

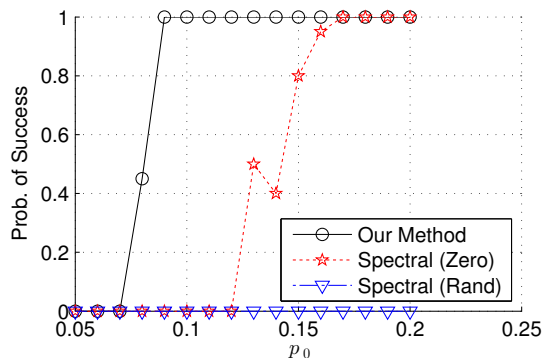


Figure 6: Comparison of our method and spectral clustering under different observation probabilities p_0 , with $n = 2000$, $K_{\min} = 200$ and $\tau = 0.1$. For spectral clustering, two imputation schemes are considered: (a) Spectral (Zero), where the missing entries are imputed with zeros, and (b) Spectral (Rand), where they are imputed with 0/1 random variables with symmetric probabilities. The result shows that our method recovers the underlying clusters with fewer observations.

this setting. The effectiveness of the proposed method and the scaling of the theoretical results are validated by simulation studies.

This work is motivated by graph clustering applications where obtaining similarity data is expensive and it is desirable to use as few observations as possible. As such, potential directions for future work include considering different sampling schemes such as active sampling, as well as dealing with sparse graphs with very few connections.

Acknowledgments

S. Sanghavi would like to acknowledge DTRA grant HDTRA1-13-1-0024 and NSF grants 1302435, 1320175 and 0954059. H. Xu is partially supported by the Ministry of Education of Singapore

through AcRF Tier Two grant R-265-000-443-112. The authors are grateful to the anonymous reviewers for their thorough reviews of this work and valuable suggestions on improving the manuscript.

Appendix A. Technical Lemmas

In this section, we provide several auxiliary lemmas required in the proof of Theorem 4. We will make use of the non-commutative Bernstein inequality. The following version is given by Tropp (2012).

Lemma 7 (Tropp, 2012) *Consider a finite sequence $\{\mathbf{M}_i\}$ of independent, random $n \times n$ matrices that satisfy the assumption $\mathbb{E}\mathbf{M}_i = 0$ and $\|\mathbf{M}_i\| \leq D$ almost surely. Let*

$$\sigma^2 = \max \left\{ \left\| \sum_i \mathbb{E} \left[\mathbf{M}_i \mathbf{M}_i^\top \right] \right\|, \left\| \sum_i \mathbb{E} \left[\mathbf{M}_i^\top \mathbf{M}_i \right] \right\| \right\}.$$

Then for all $\theta > 0$ we have

$$\begin{aligned} \mathbb{P} \left[\left\| \sum \mathbf{M}_i \right\| \geq \theta \right] &\leq 2n \exp \left(-\frac{\theta^2}{2\sigma^2 + 2D\theta/3} \right). \\ &\leq \begin{cases} 2n \exp \left(-\frac{3\theta^2}{8\sigma^2} \right), & \text{for } \theta \leq \frac{\sigma^2}{D}; \\ 2n \exp \left(-\frac{3\theta}{8D} \right), & \text{for } \theta \geq \frac{\sigma^2}{D}. \end{cases} \end{aligned} \tag{7}$$

Remark 8 *When $n = 1$, this becomes the standard two-sided Bernstein inequality.*

We will also make use of the following estimate, which follows from the structure of \mathbf{U} .

$$\left\| \mathcal{P}_{\mathcal{T}}(e_i e_j^\top) \right\|_F^2 = \|\mathbf{U}\mathbf{U}^\top e_i\|^2 + \|\mathbf{U}\mathbf{U}^\top e_j\|^2 - \|\mathbf{U}\mathbf{U}^\top e_i\|^2 \|\mathbf{U}\mathbf{U}^\top e_j\|^2 \leq \frac{2n}{K_{\min}^2}, \quad \forall 1 \leq i, j \leq n.$$

The first auxiliary lemma controls the operator norm of certain random operators. A similar result was given in Candès et al. (2011, Theorem 4.1). Our proof is different from theirs.

Lemma 9 *Suppose Ω_0 is a set of entries obeying $\Omega_0 \sim \text{Ber}_1(\rho)$. Consider the operator $\mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}}\mathcal{R}_{\Omega_0}\mathcal{P}_{\mathcal{T}}$. For some constant $C_0 > 0$, we have*

$$\|\mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}}\mathcal{R}_{\Omega_0}\mathcal{P}_{\mathcal{T}}\| < \epsilon_1$$

with high probability provided that $\rho \geq C_0 \frac{n \log n}{\epsilon_1^2 K_{\min}^2}$ and $\epsilon_1 \leq 1$.

Proof For each (i, j) , define the indicator random variable $\delta_{ij} = \mathbf{1}_{\{(i,j) \in \Omega_0\}}$. We observe that for any matrix $\mathbf{M} \in \mathcal{T}$,

$$\begin{aligned} (\mathcal{P}_{\mathcal{T}}\mathcal{R}_{\Omega_0}\mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}})\mathbf{M} &= \sum_{1 \leq i < j \leq n} \mathcal{S}_{ij}(\mathbf{M}) \\ &\triangleq \sum_{1 \leq i < j \leq n} (\rho^{-1}\delta_{ij} - 1) \left\langle \mathcal{P}_{\mathcal{T}}(e_i e_j^\top), \mathbf{M} \right\rangle \mathcal{P}_{\mathcal{T}}(e_i e_j^\top + e_j e_i^\top). \end{aligned}$$

Here $\mathcal{S}_{ij} : \mathbb{R}^{n \times n} \mapsto \mathbb{R}^{n \times n}$ is a linear self-adjoint operator with $\mathbb{E}[\mathcal{S}_{ij}] = 0$. Using the fact that $\mathcal{P}_{\mathcal{T}}(e_i e_j^\top) = (\mathcal{P}_{\mathcal{T}}(e_j e_i^\top))^\top$ and \mathbf{M} is symmetric, we obtain the bounds

$$\begin{aligned} \|\mathcal{S}_{ij}\| &\leq \rho^{-1} \left\| \mathcal{P}_{\mathcal{T}}(e_i e_j^\top) \right\|_F \left\| \mathcal{P}_{\mathcal{T}}(e_i e_j^\top + e_j e_i^\top) \right\|_F \\ &\leq \rho^{-1} \cdot 2 \left\| \mathcal{P}_{\mathcal{T}}(e_i e_j^\top) \right\|_F^2 \leq \frac{4n}{K_{\min}^2 \rho}, \end{aligned}$$

and

$$\begin{aligned} &\left\| \mathbb{E} \left[\sum_{1 \leq i < j \leq n} \mathcal{S}_{ij}^2(\mathbf{M}) \right] \right\|_F \\ &= \left\| \sum_{1 \leq i < j \leq n} \mathbb{E} \left[(\rho^{-1} \delta_{ij}^{(k)} - 1)^2 \right] \left\langle \mathcal{P}_{\mathcal{T}}(e_i e_j^\top), \mathbf{M} \right\rangle \left\langle \mathcal{P}_{\mathcal{T}}(e_i e_j^\top + e_j e_i^\top), e_i e_j^\top \right\rangle \mathcal{P}_{\mathcal{T}}(e_i e_j^\top + e_j e_i^\top) \right\|_F \\ &= (\rho^{-1} - 1) \left\| \sum_{1 \leq i < j \leq n} 2 \left\| \mathcal{P}_{\mathcal{T}}(e_i e_j^\top) \right\|_F^2 m_{i,j} \mathcal{P}_{\mathcal{T}}(e_i e_j^\top + e_j e_i^\top) \right\|_F \\ &\leq (\rho^{-1} - 1) \left\| \sum_{1 \leq i < j \leq n} 2 \left\| \mathcal{P}_{\mathcal{T}}(e_i e_j^\top) \right\|_F^2 m_{i,j} (e_i e_j^\top + e_j e_i^\top) \right\|_F \\ &\leq (\rho^{-1} - 1) \frac{4n}{K_{\min}^2} \left\| \sum_{1 \leq i < j \leq n} m_{i,j} (e_i e_j^\top + e_j e_i^\top) \right\|_F \\ &= (\rho^{-1} - 1) \frac{4n}{K_{\min}^2} \|\mathbf{M}\|_F, \end{aligned}$$

which means $\left\| \mathbb{E} \left[\sum_{1 \leq i < j \leq n} \mathcal{S}_{ij}^2 \right] \right\| \leq \frac{4n}{K_{\min}^2 \rho}$. Applying the first inequality in the Bernstein inequality (7) gives

$$\mathbb{P} \left[\left\| \sum_{1 \leq i < j \leq n} \mathcal{S}_{ij} \right\| \geq \epsilon_1 \right] \leq 2n^{2-2\beta}$$

provided $\rho \geq \frac{64\beta n \log n}{3K_{\min}^2 \epsilon_1^2}$ and $\epsilon_1 < 1$. ■

The next lemma bounds the spectral norm of certain symmetric random matrices. A related result for non-symmetric matrices appeared in Candès and Recht (2009, Theorem 6.3).

Lemma 10 *Suppose Ω_0 is a set of entries obeying $\Omega_0 \sim \text{Ber}_1(\rho)$, and \mathbf{M} is a fixed $n \times n$ symmetric matrix. Then for some constant $C_0 > 0$, we have*

$$\|(\mathcal{I} - \mathcal{R}_{\Omega_0})\mathbf{M}\| < \sqrt{C_0 \frac{n \log n}{\rho}} \|\mathbf{M}\|_\infty,$$

with high probability provided that $\rho \geq C_0 \frac{\log n}{n}$.

Proof Define δ_{ij} as before. Notice that

$$\mathcal{R}_{\Omega_0}(\mathbf{M}) - \mathbf{M} = \sum_{i < j} S_{ij} \triangleq \sum_{i < j} (\rho^{-1} \delta_{ij} - 1) m_{i,j} \left(e_i e_j^\top + e_j e_i^\top \right).$$

Here the symmetric matrix $S_{ij} \in \mathbb{R}^{n \times n}$ satisfies $\mathbb{E}[S_{ij}] = 0$, $\|S_{ij}\| \leq 2\rho^{-1} \|\mathbf{M}\|_\infty$ and the bound

$$\begin{aligned} \left\| \mathbb{E} \left[\sum_{i < j} S_{ij}^2 \right] \right\| &= (\rho^{-1} - 1) \left\| \sum_{i < j} m_{i,j}^2 \left(e_i e_i^\top + e_j e_j^\top \right) \right\| \\ &\leq (\rho^{-1} - 1) \left\| \text{diag} \left(\sum_j m_{1,j}^2, \dots, \sum_j m_{n,j}^2 \right) \right\| \\ &\leq (\rho^{-1} - 1) n \|\mathbf{M}\|_\infty^2 \leq 2\rho^{-1} n \|\mathbf{M}\|_\infty^2. \end{aligned}$$

When $\rho \geq \frac{16\beta \log n}{3n}$, we apply the first inequality in the Bernstein inequality (7) to obtain

$$\mathbb{P} \left[\left\| \sum_{i < j} S_{ij} \right\| \geq \sqrt{\frac{16\beta n \log n}{3\rho}} \|\mathbf{M}\|_\infty \right] \leq 2n \exp \left(-\frac{3 \cdot \frac{16\beta n \log n}{3\rho} \|\mathbf{M}\|_\infty^2}{8 \cdot \frac{2n}{\rho} \|\mathbf{M}\|_\infty^2} \right) \leq 2n^{1-\beta}.$$

The conclusion follows by choosing a sufficiently large constant β . ■

The third lemma bounds the infinity norm of certain random symmetric matrices. A related result is given in Candès et al. (2011, Lemma 3.1).

Lemma 11 *Suppose Ω_0 is a set of entries obeying $\Omega_0 \sim \text{Ber}_1(\rho)$, and $\mathbf{M} \in \mathcal{T}$ is a fixed symmetric $n \times n$ matrix. Then for some constant $C_0 > 0$, we have*

$$\|(\mathcal{P}_{\mathcal{T}} - \mathcal{P}_{\mathcal{T}} \mathcal{R}_{\Omega_0} \mathcal{P}_{\mathcal{T}}) \mathbf{M}\|_\infty < \epsilon_3 \|\mathbf{M}\|_\infty,$$

with high probability provided that $\rho \geq C_0 \frac{n \log n}{\epsilon_3^2 K_{\min}^2}$ and $\epsilon_3 \leq 1$.

Proof Define δ_{ij} as before. Fix an entry index (a, b) . Notice that

$$(\mathcal{P}_{\mathcal{T}} \mathcal{R}_{\Omega_0} \mathcal{P}_{\mathcal{T}} \mathbf{M} - \mathcal{P}_{\mathcal{T}} \mathbf{M})_{a,b} = \sum_{i < j} \xi_{ij} \triangleq \sum_{i < j} \left\langle (\rho^{-1} \delta_{ij}^{(k)} - 1) m_{i,j} \mathcal{P}_{\mathcal{T}} \left(e_i e_j^\top + e_j e_i^\top \right), e_a e_b^\top \right\rangle.$$

The random variable ξ_{ij} satisfies $\mathbb{E}[\xi_{ij}] = 0$ and obeys the bounds

$$|\xi_{ij}| \leq 2\rho^{-1} \left\| \mathcal{P}_{\mathcal{T}}(e_i e_j^\top) \right\|_F \left\| \mathcal{P}_{\mathcal{T}}(e_a e_b^\top) \right\|_F |m_{i,j}| \leq \frac{4n}{K_{\min}^2 \rho} \|\mathbf{M}\|_\infty$$

and

$$\begin{aligned}
 \left| \mathbb{E} \left[\sum_{i < j} \xi_{ij}^2 \right] \right| &= \left| \sum_{i < j} \mathbb{E} \left[(\rho^{-1} \delta_{ij}^{(k)} - 1)^2 \right] m_{i,j}^2 \left\langle \mathcal{P}_{\mathcal{T}} \left(e_i e_j^\top + e_j e_i^\top \right), e_a e_b^\top \right\rangle^2 \right| \\
 &\leq (\rho^{-1} - 1) \|\mathbf{M}\|_\infty^2 \sum_{i < j} \left\langle e_i e_j^\top + e_j e_i^\top, \mathcal{P}_{\mathcal{T}}(e_a e_b^\top) \right\rangle^2 \\
 &\leq 2 (\rho^{-1} - 1) \|\mathbf{M}\|_\infty^2 \left\| \mathcal{P}_{\mathcal{T}}(e_a e_b^\top) \right\|_F^2 \\
 &\leq 2 (\rho^{-1} - 1) \frac{2n}{K_{\min}^2} \|\mathbf{M}\|_\infty^2 \\
 &\leq \frac{4n}{K_{\min}^2 \rho} \|\mathbf{M}\|_\infty^2.
 \end{aligned}$$

When $\rho \geq \frac{64\beta n \log n}{3K_{\min}^2 \epsilon_3^2}$ and $\epsilon_3 \leq 1$, we apply the first inequality in the Bernstein inequality (7) with $n = 1$ to obtain

$$\mathbb{P} \left[\left| (\mathcal{P}_{\mathcal{T}} \mathcal{R}_{\Omega_0} \mathcal{P}_{\mathcal{T}} \mathbf{M} - \mathcal{P}_{\mathcal{T}} \mathbf{M})_{a,b} \right| \geq \epsilon_3 \|\mathbf{M}\|_\infty \right] \leq 2 \exp \left(- \frac{3\epsilon_3^2 \|\mathbf{M}\|_\infty^2}{8 \frac{4n}{K_{\min}^2 \rho} \|\mathbf{M}\|_\infty^2} \right) \leq 2n^{-2\beta}.$$

Applying the union bound then yields

$$\mathbb{P} \left[\|\mathcal{P}_{\mathcal{T}} \mathcal{R}_{\Omega_0} \mathcal{P}_{\mathcal{T}} \mathbf{M} - \mathcal{P}_{\mathcal{T}} \mathbf{M}\|_\infty \geq \epsilon_3 \|\mathbf{M}\|_\infty \right] \leq 2n^{2-2\beta}.$$

■

The last lemma bounds the matrix infinity norm of $\mathcal{P}_{\mathcal{T}} \mathcal{P}_{\Omega}(\mathbf{S})$ for a ± 1 random matrix \mathbf{S} .

Lemma 12 *Suppose $\Omega \sim \text{Ber}_0 \left(\frac{2p_0\tau}{1-p_0+2p_0\tau} \right)$ and $\mathbf{S} \in \mathbb{R}^{n \times n}$ has i.i.d. symmetric ± 1 entries. Under the assumption of Theorem 4, for some constant C_0 , we have with high probability*

$$\|\mathcal{P}_{\mathcal{T}} \mathcal{P}_{\Omega}(\mathbf{S})\|_\infty \leq C_0 \sqrt{\frac{p_0 n \log n}{K_{\min}^2}}.$$

Proof By triangle inequality, we have

$$\|\mathcal{P}_{\mathcal{T}} \mathcal{P}_{\Omega}(\mathbf{S})\|_\infty \leq \|\mathbf{U} \mathbf{U}^T \mathcal{P}_{\Omega}(\mathbf{S})\|_\infty + \|\mathcal{P}_{\Omega}(\mathbf{S}) \mathbf{U} \mathbf{U}^T\|_\infty + \|\mathbf{U} \mathbf{U}^T \mathcal{P}_{\Omega}(\mathbf{S}) \mathbf{U} \mathbf{U}^T\|_\infty,$$

so it suffices to show that each of these three terms are bounded by $C \sqrt{\frac{p_0 n \log n}{K_{\min}^2}}$ w.h.p. for some constant C . Under the assumption on Ω and \mathbf{S} in the lemma statement, each pair of symmetric entries of $\mathcal{P}_{\Omega}(\mathbf{S})$ equals ± 1 with probability $\rho \triangleq \frac{p_0\tau}{1-p_0+2p_0\tau}$ and 0 otherwise; notice that $\rho \leq \frac{p_0}{2}$ since $\tau \leq \frac{1}{2}$. Let $(s^{(i)})^T$ be the i th row of $\mathbf{U} \mathbf{U}^T$. From the structure of \mathbf{U} , we know that for all i and j ,

$$\left| s_j^{(i)} \right| \leq \frac{1}{K_{\min}},$$

and for all i ,

$$\sum_{j=1}^n \left(s_j^{(i)}\right)^2 \leq \frac{1}{K_{\min}}.$$

We now bound $\|\mathbf{U}\mathbf{U}^T\mathcal{P}_\Omega(\mathbf{S})\|_\infty$. For simplicity, we focus on the $(1, 1)$ entry of $\mathbf{U}\mathbf{U}^T\mathcal{P}_\Omega(\mathbf{S})$ and denote this random variable as X . We may write X as $X = \sum_i s_i^{(1)} (\mathcal{P}_\Omega(\mathbf{S}))_{i,1}$, for which we have

$$\begin{aligned} \mathbb{E} \left[s_i^{(1)} (\mathcal{P}_\Omega(\mathbf{S}))_{i,1} \right] &= 0, \\ \left| s_i^{(1)} (\mathcal{P}_\Omega(\mathbf{S}))_{i,1} \right| &\leq \left| s_i^{(1)} \right| \leq \frac{1}{K_{\min}}, \quad a.s. \\ \text{Var}(X) &= \sum_{i:(i,1) \in \Omega} (s_i^{(1)})^2 \cdot 2\rho \leq \frac{p_0}{K_{\min}}. \end{aligned}$$

Applying the standard Bernstein inequality then gives

$$\mathbb{P} \left[|X| > C \sqrt{\frac{p_0 n \log n}{K_{\min}^2}} \right] \leq 2 \exp \left[- \left(C^2 \frac{p_0 n \log n}{K_{\min}^2} \right) / \left(2 \frac{p_0}{K_{\min}} + \frac{2C \sqrt{p_0 n \log n}}{3K_{\min}^2} \right) \right].$$

Under the assumption of Theorem 4, the right hand side above is bounded by $2n^{-12}$. It follows from the union bound that $\|\mathbf{U}\mathbf{U}^T\mathcal{P}_\Omega(\mathbf{S})\|_\infty \leq C \sqrt{\frac{p_0 n \log n}{K_{\min}^2}}$ w.h.p. Clearly, the same bound holds for $\|\mathcal{P}_\Omega(\mathbf{S})\mathbf{U}\mathbf{U}^T\|_\infty$. Finally, let K be the size of the cluster that node j is in. Observe that due to the structure of $\mathbf{U}\mathbf{U}^T$, we have

$$\left(\mathbf{U}\mathbf{U}^T\mathcal{P}_\Omega(\mathbf{S})\mathbf{U}\mathbf{U}^T\right)_{i,j} = \sum_l \left(\mathbf{U}\mathbf{U}^T\mathcal{P}_\Omega(\mathbf{S})\right)_{i,l} \left(\mathbf{U}\mathbf{U}^T\right)_{l,j} \leq \frac{1}{K} \cdot K \cdot \left\| \mathbf{U}\mathbf{U}^T\mathcal{P}_\Omega(\mathbf{S}) \right\|_\infty,$$

which implies $\|\mathbf{U}\mathbf{U}^T\mathcal{P}_\Omega(\mathbf{S})\mathbf{U}\mathbf{U}^T\|_\infty \leq \|\mathbf{U}\mathbf{U}^T\mathcal{P}_\Omega(\mathbf{S})\|_\infty$. This completes the proof of the lemma. \blacksquare

Appendix B. Proof of Theorem 5

We use a standard information theoretical argument, which improves upon a related proof by Chaudhuri et al. (2012). Let K be the size of the clusters (which are assumed to have equal size). For simplicity we assume n/K is an integer. Let \mathcal{F} be the set of all possible partition of n nodes into n/K clusters of equal size K . Using Stirling's approximation, we have

$$M \triangleq |\mathcal{F}| = \frac{1}{(n/K)!} \binom{n}{K} \binom{n-K}{K} \cdots \binom{K}{K} \geq \left(\frac{n}{3K}\right)^{n(1-\frac{1}{K})} \geq c_1^{\frac{1}{2}n},$$

which holds for $K = \Theta(n)$.

Suppose the clustering \mathbf{Y} is chosen uniformly at random from \mathcal{F} , and the graph \mathbf{A} is generated from \mathbf{Y} according to the planted partition model with partial observations, where we use $a_{ij} = ?$ for

unobserved pairs. We use $\mathbb{P}_{\mathbf{A}|\mathbf{Y}}$ to denote the distribution of \mathbf{A} given \mathbf{Y} . Let $\hat{\mathbf{Y}}$ be any measurable function of the observation \mathbf{A} . A standard application of Fano's inequality and the convexity of the mutual information (Yang and Barron, 1999) gives

$$\sup_{\mathbf{Y} \in \mathcal{F}} \mathbb{P} \left[\hat{\mathbf{Y}} \neq \mathbf{Y} | \mathbf{Y} \right] \geq 1 - \frac{M^{-2} \sum_{\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)} \in \mathcal{F}} D \left(\mathbb{P}_{\mathbf{A}|\mathbf{Y}^{(1)}} \| \mathbb{P}_{\mathbf{A}|\mathbf{Y}^{(2)}} \right) + \log 2}{\log M}, \quad (8)$$

where $D(\cdot|\cdot)$ denotes the KL-divergence. We now upper bound this divergence. Given $\mathbf{Y}^{(l)}$, $l = 1, 2$, the $a_{i,j}$'s are independent of each other, so we have

$$D \left(\mathbb{P}_{\mathbf{A}|\mathbf{Y}^{(1)}} \| \mathbb{P}_{\mathbf{A}|\mathbf{Y}^{(2)}} \right) = \sum_{i,j} D \left(\mathbb{P}_{a_{i,j}|\mathbf{Y}^{(1)}} \| \mathbb{P}_{a_{i,j}|\mathbf{Y}^{(2)}} \right).$$

For each pair (i, j) , the KL-divergence is zero if $y_{i,j}^{(1)} = y_{i,j}^{(2)}$, and otherwise satisfies

$$\begin{aligned} D \left(\mathbb{P}_{a_{i,j}|\mathbf{Y}^{(1)}} \| \mathbb{P}_{a_{i,j}|\mathbf{Y}^{(2)}} \right) &\leq p_0(1 - \tau) \log \frac{p_0(1 - \tau)}{p_0\tau} + p_0\tau \log \frac{p_0\tau}{p_0(1 - \tau)} + (1 - p_0) \log \frac{1 - p_0}{1 - p_0} \\ &= p_0(1 - 2\tau) \log \frac{1 - \tau}{\tau} \\ &\leq p_0(1 - 2\tau) \left(\frac{1 - \tau}{\tau} - 1 \right) \\ &\leq c_2 p_0(1 - 2\tau)^2, \end{aligned}$$

where $c_2 > 0$ is a universal constant and the last inequality holds under the assumption $\tau > 1/100$. Let N be the number of pairs (i, j) such that $y_{i,j}^{(1)} \neq y_{i,j}^{(2)}$. When $K = \Theta(n)$, we have

$$N \leq |\{(i, j) : y_{i,j}^{(1)} = 1\} \cup \{(i, j) : y_{i,j}^{(2)} = 1\}| \leq n^2.$$

It follows that $D \left(\mathbb{P}_{\mathbf{A}|\mathbf{Y}^{(1)}} \| \mathbb{P}_{\mathbf{A}|\mathbf{Y}^{(2)}} \right) \leq N \cdot c_2 p_0(1 - 2\tau)^2 \leq c_2 n^2 p_0(1 - 2\tau)^2$. Combining pieces, for the left hand side of (8) to be less than $1/4$, we must have $p_0(1 - 2\tau)^2 \geq C \frac{1}{n}$.

References

- N. Alon, M. Krivelevich, and B. Sudakov. Finding a large hidden clique in a random graph. In *Proceedings of the 9th annual ACM-SIAM Symposium on Discrete Algorithms*, pages 457–466, 1998.
- B. Ames and S. Vavasis. Nuclear norm minimization for the planted clique and biclique problems. *Mathematical Programming*, 129(1):69–89, 2011.
- M. F. Balcan and P. Gupta. Robust hierarchical clustering. In *Proceedings of the Conference on Learning Theory (COLT)*, 2010.
- N. Bansal, A. Blum, and S. Chawla. Correlation clustering. In *Proceedings of the 43rd Symposium on Foundations of Computer Science*, 2002.

- B. Bollobás and A. D. Scott. Max cut for random graphs with a planted partition. *Combinatorics, Probability and Computing*, 13(4-5):451–474, 2004.
- R. B. Boppana. Eigenvalues and graph bisection: An average-case analysis. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science*, pages 280–285, 1987.
- E. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- E. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):11, 2011.
- T. Carson and R. Impagliazzo. Hill-climbing finds random planted bisections. In *Proceedings of the 12th annual ACM-SIAM Symposium on Discrete Algorithms*, pages 903–909, 2001.
- V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2):572–596, 2011.
- M. Charikar, V. Guruswami, and A. Wirth. Clustering with qualitative information. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 524–533, 2003.
- K. Chaudhuri, F. Chung, and A. Tsiatas. Spectral clustering of graphs with general degrees in the extended planted partition model. *Journal of Machine Learning Research*, 2012:1–23, 2012.
- Y. Chen, A. Jalali, S. Sanghavi, and C. Caramanis. Low-rank matrix recovery from errors and erasures. *IEEE Transactions on Information Theory*, 59(7):4324–4337, 2013.
- A. Condon and R. M. Karp. Algorithms for graph partitioning on the planted partition model. *Random Structures and Algorithms*, 18(2):116–140, 2001.
- A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, 2011.
- E. D. Demaine and N. Immorlica. Correlation clustering with partial information. *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques*, pages 71–80, 2003.
- E. D. Demaine, D. Emanuel, A. Fiat, and N. Immorlica. Correlation clustering in general weighted graphs. *Theoretical Computer Science*, 361(2):172–187, 2006.
- D. Emmanuel and A. Fiat. Correlation clustering minimizing disagreements on arbitrary weighted graphs. In *Proceedings of the 11th Annual European Symposium on Algorithms*, pages 208–220, 2003.
- B. Eriksson, G. Dasarathy, A. Singh, and R. Nowak. Active clustering: Robust and efficient hierarchical clustering using adaptively selected similarities. *Arxiv preprint arXiv:1102.3887*, 2011.
- M. Ester, H. Kriegel, and X. Xu. A database interface for clustering in large spatial databases. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining*, pages 94–99, 1995.

- M. Fazel. *Matrix Rank Minimization with Applications*. PhD thesis, Stanford University, 2002.
- U. Feige and J. Kilian. Heuristics for semirandom graph problems. *Journal of Computer and System Sciences*, 63(4):639–671, 2001.
- J. Giesen and D. Mitsche. Reconstructing many partitions using spectral techniques. In *Fundamentals of Computation Theory*, pages 433–444. Springer, 2005.
- D. Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.
- P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: Some first steps. *Social networks*, 5(2):109–137, 1983.
- D. Hsu, S. M. Kakade, and T. Zhang. Robust matrix decomposition with sparse corruptions. *IEEE Transactions on Information Theory*, 57(11):7221–7234, 2011.
- B. Hunter and T. Strohmer. Spectral clustering with compressed, incomplete and inaccurate measurements. Available at <https://www.math.ucdavis.edu/~strohmer/papers/2010/SpectralClustering.pdf>, 2010.
- M. Jerrum and G. B. Sorkin. The metropolis algorithm for graph bisection. *Discrete Applied Mathematics*, 82(1-3):155–175, 1998.
- B. W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell System Technical Journal*, 49(2):291–307, 1970.
- A. Krishnamurthy, S. Balakrishnan, M. Xu, and A. Singh. Efficient active algorithms for hierarchical clustering. In *Proceedings of the 29th International Conference on Machine Learning*, pages 887–894, 2012.
- X. Li. Compressed sensing and matrix completion with constant proportion of corruptions. *Constructive Approximation*, 37(1):73–99, 2013.
- Z. Lin, M. Chen, L. Wu, and Y. Ma. The Augmented Lagrange Multiplier Method for Exact Recovery of Corrupted Low-Rank Matrices. *UIUC Technical Report UILU-ENG-09-2215*, 2009.
- C. Mathieu and W. Schudy. Correlation clustering with noisy input. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 712–728, 2010.
- F. McSherry. Spectral partitioning of random graphs. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 529–537, 2001.
- N. Mishra, I. Stanton R. Schreiber, and R. E. Tarjan. Clustering social networks. In *Algorithms and Models for Web-Graph*, pages 56–67. Springer, 2007.
- S. Oymak and B. Hassibi. Finding dense clusters via low rank + sparse decomposition. Available on arXiv:1104.5186v1, 2011.
- K. Rohe, S. Chatterjee, and B. Yu. Spectral clustering and the high-dimensional stochastic block-model. *The Annals of Statistics*, 39(4):1878–1915, 2011.

- O. Shamir and N. Tishby. Spectral clustering on a budget. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, pages 661–669, 2011.
- R. Shamir and D. Tsur. Improved algorithms for the random cluster graph model. *Random Structures & Algorithms*, 31(4):418–449, 2007.
- C. Swamy. Correlation clustering: maximizing agreements via semidefinite programming. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2004.
- J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.
- R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *Arxiv preprint arxiv:1011.3027*, 2010.
- K. Voevodski, M. F. Balcan, H. Roglin, S. H. Teng, and Y. Xia. Efficient clustering with limited distance information. *arXiv preprint arXiv:1009.5168*, 2010.
- U. von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.
- Yahoo!-Inc. Graph partitioning. Available at <http://research.yahoo.com/project/2368>, 2009.
- Y. Yang and A. Barron. Information-theoretic determination of minimax rates of convergence. *The Annals of Statistics*, 27(5):1564–1599, 1999.