Electronic Colloquium on Computational Complexity, Report No. 125 (2009)

Connections Between Unique Games and Multicut *

David Steurer† Nisheeth K. Vishnoi‡

Abstract

In this paper we demonstrate a close connection between Unique Games and MultiCut. In MultiCut, one is given a *network graph* and a *demand graph* on the same vertex set and the goal is to remove as few edges from the network graph as possible such that every two vertices connected by a demand edge are separated. On the other hand, Unique Games is a certain family of constraint satisfaction problems.

In one direction, we show that, at least with respect to current algorithmic techniques, MultiCut is not harder than Unique Games. Specifically, we can adapt most known algorithms for Unique Games to work for MultiCut and obtain new approximation guarantees for MultiCut that depend on the maximum degree of the demand graph. This degree plays the same role as the alphabet size plays in approximation guarantees for Unique Games.

In the other direction, we show that MultiCut is not easier than Unique Games (Γ Max 2 Lin to be precise). We exhibit a reduction from Unique Games to MultiCut so that the fraction of edges in the optimal multicut is up to a factor of 2 equal to the fraction of constraint violated by the optimal assignment for the Unique Games instance. In contrast to the vast majority of Unique Games reductions whose analysis relies on Fourier analysis and isoperimetric inequalities, this reduction is simple and the analysis is elementary. Further, the maximum degree of the demand graph in the instance produced by the reduction is less than the size of the alphabet size in the Unique Games instance.

Our results rely on a simple but previously unknown characterization of MultiCut in terms of L_1 metrics.

*Part of this work was done when the two authors were at Laboratoire de Recherche en Informatique, Orsay, France in December 2008.

†Princeton Univeristy. Email: dsteurer@cs.princeton.edu
‡Microsoft Research, Bangalore, India. Email: nisheeth.vishnoi@gmail.com.

ISSN 1433-8092

1 Introduction

Minimum Multicut. An instance of Minimum Multicut (MultiCut) is specified by two graphs G and H on the same vertex set. We refer to G as the *network graph* and to H as the *demand graph*. The goal is to find a partition of the vertex set such that all demand edges are separated with the objective to minimize the fraction of network edges cut.

Unique Games. Unique Games is a constraint satisfaction problem where, given a constraint graph $G(V, E)$, a label set $[k]$ and for each edge $e = (u, v)$, a bijective constraint $\pi_{uv} : [k] \mapsto [k]$, the goal is to assign to each vertex in G a label from $[k]$ so as to maximize the fraction of the constraints that are *satisfied*. A constraint $e = (u, v)$ is said to be *satisfied* by an assignment if u is assigned a label i and v is assigned a label j such that $\pi_{uv}(i) = j$. over all labelings. The Unique Games Conjecture (UGC) of Khot [Kho02] asserts that for such a constraint satisfaction problem, for arbitrarily small constants $\eta, \zeta > 0$, it is NP-hard to decide whether there is a labeling that satisfies $1 - \eta$ fraction of the constraints or, for every labeling, the fraction of the constraints satisfied is at most ζ as long as the size of the label set, k, is allowed to grow as a function of η and ζ.

1.1 Previous and Related Work

MultiCut. The MultiCut problem is a generalization of the classic minimum (s, t)-cut problem and has received a considerable amount of attention (see the book by Vazirani [Vaz01] for a comprehensive treatment of this problem). The MultiCut problem is NP-hard, and the best approximation algorithm for this problem is due to Garg, Vazirani and Yannakakis [GVY96] who build on the *region growing* technique of Leighton and Rao [LR99] to give a $O(\log |E(H)|) = O(\log n)$ factor approximation algorithm for this problem. The result of Garg, Vazirani and Yannakakis relies on a certain natural linear programming (LP) relaxation of the MultiCut problem. This relaxation is inspired by the LP formulation of the MaxFlow-MinCut problem. They also prove that, unlike in the case of the MaxFlow-MinCut problem, there is an integrality gap of $\Omega(\log n)$ for the LP that they consider. Subsequently there have been attempts, for instance by Agarwal, Charikar, Makarychev and Makarychev [ACMM05], to reduce the gap of $\Theta(\log n)$ to $o(\log n)$ using a semi-definite programming (SDP) based approach. The hope was to strengthen the Garg-Vazirani-Yannakakis relaxation for MultiCut in the world of SDP in a manner similar to how Arora, Rao and Vazirani [ARV04], for the Sparsest Cut problem, strengthened the Leighton-Rao LP relaxation to a SDP and obtained a $O(\sqrt{\log n})$ approximation algorithm for the Sparsest Cut problem. In such an attempt, Agarwal *et al.* proved that the gap remains $\Omega(\log n)$ even when one goes far beyond SDPs: to the computationally infeasible world of optimizing over the space of L_1 embeddable metrics. This is in stark contrast with the Sparsest Cut problem where it was discovered by Aumann and Rabani [AR98] and Linial, London and Rabinovich [LLR95] that the L_1 tightening of the LP considered by Leighton and Rao for the Sparsest Cut problem has no integrality gap.

On the hardness side, it is known that the problem is as hard to approximate as the Vertex Cover problem, even when H is a tree. The current best unconditional result for Vertex Cover is about 1.36 due to Dinur and Safra [DS05]. Then there are conditional results due to Chawla *et al.* [CKK+05] and Khot and Vishnoi [KV05] who prove that, assuming the Unique Games Conjecture (UGC) of Khot [Kho02], there is no constant factor approximation to the MultiCut problem.

Unique Games. Since its origin, the UGC has been successfully used to prove (often optimal) hardness of approximation results for several important NP-hard problems such as Vertex Cover [KR03], Maximum Cut [KKMO04] and a wide class of constraint satisfaction problems [Rag08]. The list of problems for which the best inapproximability result rests on the validity of the UGC is growing rapidly and in a relatively short time since the conjecture was made, it has become one of the key open problems in complexity theory and approximation algorithms. The validity of this conjecture remains under scrutiny.

Several attempts have been made to disprove the UGC. Starting with the original paper of Khot [Kho02] itself, to the papers by Trevisan [Tre05], Gupta and Talwar [GT06], Charikar, Makarychev and Makarychev [CMM06a] and Chlamtac, Makarychev and Makarychev [CMM06b]. Needless to say, they all fall short of disproving the UGC. Except the result of Gupta and Talwar, which is a LP based approach, the other results are based on SDP relaxations of the Unique Games problem. Roughly, they all lead to the conclusion that the UGC cannot be strengthened to the case when η, ζ to very small $o(1)$ terms.

On attempts to prove the UGC, very recently, an approach to put it on a firm ground has been ruled out by a result of Raz [Raz08] and suitably generalized by Barak *et al.* [BHH+08]. The latter result proves that one cannot hope to amplify hardness of a UGC instance by taking tensor products, a standard approach to amplify hardness.

1.2 Main Results

We show that there is a close connection between the MultiCut problem on instances with maximum degree Δ in the demand graph and the Unique Games problem where the number of labels is about Δ.

***Hardness result for* MultiCut *assuming* Unique Games.** [1] [2] The starting point of our results is a simple reduction that, assuming the UGC, the MultiCut problem is hard to approximate to within any constant factor. The previous proofs of this theorem, following the *dictatorship test* paradigm, rely on fourier analysis. Our reduction, on the other hand, relies on a new relaxation of the MultiCut problem, and is approximation preserving in a sense we will describe later.

SDP *based algorithms for* MultiCut *from* Unique Games *algorithms.* We show that SDP rounding algorithms developed for Unique Games apply with simple modifications to the (seemingly harder) MultiCut problem with bounded degree demand graph instances. For instance an algorithm of Charikar *et al.* [CMM06a] can recover a labeling of value at least $1 - \sqrt{\varepsilon \log k}$ given an SDP solution of value $1 - \varepsilon$. The analogous result for MultiCut we show is to recover a multicut of instance at most $\sqrt{\varepsilon \log(\Delta/\varepsilon)}$ from an SDP solution of value at most ε. Here, Δ is the largest degree in the demand graph. Similarly, from a result of Chlamtac *et al.* [CMM06b] we deduce an algorithm for MultiCut which outputs a multicut of value at most $\varepsilon \sqrt{\log n \log(\Delta/\varepsilon)}$ from an SDP solution of value at most ε.

Characterization of* MultiCut *by* L_1 *metrics. Our results rely on a new tool which is a relaxation of the MultiCut problem as an optimization problem over the class of L_1 embeddable metrics. We prove that this relaxation captures the optimal value within a factor of 2. Though the optimization problem over L_1 embeddable metrics is not known to be computationally feasible, it allows us to prove the hardness result and consider SDP relaxations to it and derive several new algorithmic results for the MultiCut problem mentioned above.

1.3 Further Connections

Easiness results for tensor products of* MultiCut *instances. A typical approach to constructing hard instances is to study how the optimal value behaves under taking products. This approach failed for Unique Games as shown by Raz [Raz08] and Barak *et al.* [BHH+08]. In the spirit of translating results for Unique Games to MultiCut, we prove an analogous result: if the demand graph of a MultiCut instance is a product of t demand graphs, in each of which the maximum degree is bounded by Δ, then in the product graph, even though the maximum degree of the demand graph could be as high as Δ^t, the approximation factors just depend on Δ. As in Barak *et al.*, The key here is to show that if one is given a vector solution to a certain SDP of the MultiCut problem such that all vectors have non-negative entries, then one can get approximation factors which do not depend on the degree of the demand graph. Then one only needs to show how to translate general SDP solutions to ones with non-negative coordinates such that the cost of the non-negative solution is not too high compared to the original one.

Algorithm for* MultiCut *instance with demand graphs without large independent set. One of our algorithmic results stated before implies that if the degree of the demand graph is bounded by a constant then there is essentially a $O(\sqrt{\log n})$ approximation algorithm for the MultiCut problem. Bounded degree graphs are sparse but and will typically have independent sets of size $\Omega(n)$. We prove a somewhat complementary result: if the size of the largest independent set in the demand graph without taking into account isolated vertices is α,, then there is a $\log \alpha$ factor approximation for the MultiCut problem. Firstly, this result implies the Garg, Vazirani and Yannakakis result. Secondly, it implies that if the average degree of the demand graph is $\Omega(n)$ then there is a constant factor approximation for MultiCut. This

[1] A similar result was communicated to us by Julia Chuzhoy.

[2] We quickly observe that the Sherali–Adams SDP gap for Unique Games due to [RS09] can be translated using our reduction to give strong $\omega(1)$ integrality gap for MultiCut. This integrality gap is not implied by the integrality gap for MultiCut in [ACMM05].

result shows that, for instance, for almost all dense demand graphs (say sampled from $\mathcal{G}(n, 1/2)$), MultiCut can be approximated within a factor of $O(\log\log n)$.

1.4 Interpretations and Future Directions

The main appeal of our work is the viewpoint and the directions that seem to emerge from it. One challenge that emerges from the connections above is to come up with algorithms for Unique Games that do not generalize to MultiCut on bounded degree demand graph . The only possible candidate we are aware of is an algorithm for Unique Games on expanding constraint graphs [AKK+08]. The first obstacle in adapting this algorithm for MultiCut is to define the right notion of "expanding MultiCut instances" (expansion of the network graph doesn't seem to be enough).

Another interpretation, and the original motivation for the paper, is that our results give hope that it might be possible to reduce MultiCut on bounded degree demand graphs to Unique Games. This could be an avenue of important future research.

Yet another possibility one can explore is to try to base the hardness results relying on the UGC on the following tempting hypothesis. (Our reduction shows that the UGC implies this hypothesis.)

Multicut Hypothesis. For every small enough constant $\varepsilon > 0$, there is a $\Delta(\varepsilon)$ such that it is NP-hard to decide, given an instance of MultiCut with the maximum degree of the demand graph is bounded by $\Delta(\varepsilon)$, whether the optimal multicut is of value less than ε or more than $1-\varepsilon$.

1.5 Techniques

In this section we give an overview of our techniques.

Hardness. The hardness reduction is simple enough to be described here. It takes a Unique Games instance $\mathcal{U}$ with constraint graph $G(V, E)$, label set $[k]$, and permutations $\{\pi_e\}_{e\in E}$ and constructs the following MultiCut instance: for every vertex $v \in V$, there is a *cloud* of k vertices indexed by (v, i) where $i \in [k]$. There is a demand edge between every pair (v, i) and (v, j) for $i \neq j \in [k]$. There is a network edge between (u, i) and (v, j) if (u, v) is a constraint edge in $\mathcal{U}$ and $\pi_{uv}(i) = j$. Thus, the demand graph consists of cliques over the clouds of size k and the network edges consist of a perfect matching between the clouds corresponding to a constraint in $\mathcal{U}$. The network graph constructed is sometimes referred to as the *label extended* graph corresponding to $\mathcal{U}$. If we impose, w.l.o.g., an extra structure on the bijections in $\mathcal{U}$ that they are *linear*[3], it is easy to see that a labeling satisfying $1-\varepsilon$ fraction of the constraints gives a multicut in the label extended graph cutting at most ε fraction of the network edges. The linearity property of the constraints implies that if $\Lambda\colon V \to \{0, 1, \ldots, k-1\}$ is a labeling, then $(\Lambda + j) \mod k$, for every $j \in [k]$, satisfies exactly the same constraints as Λ. Hence, the multicut in the label extended graph of value ε consists of the following k parts: $\{(v, (\Lambda(v)+j) \mod k) : v \in V\}$ where $j \in \{0, 1, \ldots, k-1\}$.

On the other hand, if there is a multicut in the reduced graph of value at most ε, then one can find a labeling of $\mathcal{U}$ which satisfies at least $1-2\varepsilon$ fraction of the constraints. To prove this first one notices that, by definition, a multicut is a collection of independent sets in the demand graph. Hence, every part in the multicut will contain at most one candidate label for a vertex v of $\mathcal{U}$. Thus, one strategy is to sample a random part from the multicut and assign to a previously unassigned vertices the label suggested by it. Formally, if $v \in V$ has not been assigned a label yet and (v, i) appears in the part sampled from the multicut in the label extended graph, assign v label i. Keep doing this until all vertices have been assigned a label. (This will happen with probability one.) Notice that the label of two vertices u, v joined by an edge in $\mathcal{U}$ will have inconsistent labels if there is some i such that edge between (u, i) and $(v, \pi_{uv}(i))$ is cut the first time one of u or v gets its label. Now, one can show using a simple argument that the expected number of edges $uv \in E$ that get assigned inconsistent labels is at most two times the fraction of edges in the multicut we started off with.

SDP **Algorithms.** The basic SDP for MultiCut requires one to come up with a unit vector for every vertex in the graph such that the vectors corresponding to any two of them connected by a demand edge are orthogonal. The goal is then to find such a vector solution which minimizes the average squared euclidean distance between pairs of vertices

[3] See Section 2.2 for the definitions

connected by a network edge. A standard SDP for (linear) Unique Games on label set $[k]$, on the other hand, requires one to come up with a set of k orthonormal vectors for every vertex so as to maximize the average *correlation* between the two sets of orthonormal vectors connected by an edge in the constraint graph. A bit more formally, if $\{u_i\}_{i\in[k]}$ and $\{v_i\}_{i\in[k]}$ are two orthonormal set of vectors corresponding to an edge (u, v) in the Unique Games instance with the bijection π between them, then the correlation between them is defined to be $\mathbb{E}_{i\in[k]}\, u_i v_{\pi(i)}$. It is easily observed that up to a simple affine transformation of the objective function, this SDP is the same as the one for MultiCut for the instance produced by our reduction.

A typical rounding algorithm for Unique Games can be interpreted as converting the SDP solution into a distribution over cuts in the label extended graph. One such approach [CMM06a] is the following: pick a random gaussian vector g and a threshold $\tau \in \mathbb{R}_+$ and output the set of labels for a vertex v in the Unique Games instance which have projection at least τ on g. Formally, let $S_v := \{i \in [k] : \langle v_i, g\rangle \geqslant \tau\}$. The choice of τ is such that the expected cardinality of S_v is one. Since $\{v_i\}_{i\in[k]}$ is an orthonormal set, $\tau = \Theta(\sqrt{\log k})$ suffices. Further if $\frac{1}{2}\,\mathbb{E}_{i\in[k]}\,\|u_i - v_{\pi_{uv}}(i)\|^2 = \varepsilon_{uv}$ for an edge uv in the Unique Games instance, the probability that there is some i such that $i \in S_u$ but $\pi_{uv}(i) \notin S_v$ is $O(\sqrt{\varepsilon_{uv}\log k})$. Now assign a random label for every vertex from the set S_v. It is not difficult to see now that this strategy proves that the SDP value for a Unique Games instance being at least $1-\varepsilon$ implies a labeling which satisfies $1 - O(\sqrt{\varepsilon \log k})$ fraction of the constraints.

Recall that there are two ways in which the MultiCut instances produced by our reduction are special: (1) the demand graphs are union of disjoint cliques of size k (hence, maximum degree of the demand graph is $k-1$), and (2) the network edges are a union of *perfect matchings* (corresponding to the bijections between pairs of cliques in the demand graph. When we systematically study various SDP based algorithms for Unique Games (like the one described above), we observe that all these algorithms do not make critical use of the two structural properties listed above. They can be easily modified to give algorithms with similar guarantees for general MultiCut instances where the role of the the label set is assumed by the maximum degree of the demand graph.

In the case of a MultiCut instance where the maximum degree of the demand graph is Δ, we do exactly the same rounding pretending our MultiCut instance was one gotten from the hardness reduction applied to a Unique Games instance with $\Delta - 1$ labels. We pick a random gaussian g and a threshold τ and first pick all vertices in the MultiCut instance whose vectors have projection more than τ on g. Since all vectors are unit and those joined by an edge in the demand graph are orthogonal, the probability that a pair of orthogonal vectors has more than τ projection on a random gaussian is $N^2(\tau)$. Here, $N(\tau)$ is the probability that a $N(0, 1)$ random variable takes value more than τ. It could happen that two vertices s, t adjacent in the demand graph end up having projection τ or more. This probability is $N^2(\tau)$. Hence, the probability that, for a given vertex, there is some neighbor of it with this projection τ or more is at most $\Delta \cdot N^2(\tau)$. This is where the maximum degree of the demand graph comes in. We have no other choice but to create a separate component in the multicut for every demand-neighbor of a vertex with projection more than τ. We keep doing this until we have found a multicut. The rest of the calculation is almost the same as that for Unique Games. One can show that the probability that a network edge uv gets *cut* (one having projection more and the other less than τ) is at most $O(\sqrt{\varepsilon_{uv}}\tau N(\tau))$. Here ε_{uv} is the squared euclidean distance between the vectors corresponding to u and v. Hence, using an argument similar to the one use in the soundness of the hardness reduction one can show that the expected size of the multicut produced is at most $O(\sqrt{\varepsilon \log(\Delta/\varepsilon)} + \varepsilon)$. Here ε is the SDP value for the MultiCut instance. Choosing $\tau := O(\sqrt{\log(\Delta/\varepsilon)})$ we get that the size of the multicut output by this algorithm is at most $O(\sqrt{\varepsilon \log(\Delta/\varepsilon)})$.

To summarize, using the SDP solution for the MultiCut problem, we come up with a distribution over cuts of the vertex set. This distribution has two nice properties: (1) it is a distribution over independent sets in the demand graph and, (2) that the expected number of cuts that separate a network edge is related to the squared euclidean distance between the vectors corresponding to the two endpoints of the network edge. This abstraction leads us to a L_1 characterization of MultiCut within a factor of 2. (See Section 2.1.1.) The SDPs are now easily seen as relaxations for this characterization and rounding algorithms as ways to convert general SDP solutions to solutions for this L_1 program.

Once we have such a distribution over independent sets over the demand graphs, the strategy is similar to that in the soundness proof of the hardness reduction: sample a sequence $T_1, T_2, \ldots$ of independent sets from this distribution. Define $S_i := T_i \setminus \cup_{j=1}^{i-1} S_j$. Stop sampling once every vertex of the MultiCut instance is in some S_i. Output the sets $S_1, S_2, \ldots$ as the multicut.

1.6 Organization

In section 2, we give a more detailed overview over the reductions and relaxations mentioned in the introduction. In section 3, we show a representative proof. In appendix A, we prove the properties of MultiCut relaxations described in section 2. In appendix B, we give a detailed proof of the properties of the reduction from Unique Games to MultiCut.

2 Relaxations and Reductions for MultiCut

Let us recall the definition of MultiCut and introduce additional notation. A MultiCut instance is specified by two graphs G and H on the same vertex V. We refer to G as the *network graph* and to H as the *demand graph*. The goal is to find a partition P of the vertex set V such that all demand edges are separated with the objective to minimize the number of network edges that are separated. Formally,

$$\mathsf{OPT}(G,H)\colon \quad \text{minimize} \quad \mathop{\mathbb{P}}_{(u,v)\in E(G)} \left\{P(u) \neq P(v)\right\} \tag{1}$$

$$\text{subject to} \quad \forall_{(s,t)\in E(H)} \quad P(s) \neq P(t)\,. \tag{2}$$

Here, $P(u) \subseteq V$ denotes the cluster of the partition P that contains the vertex $u \in V$. We denote the fraction of network edges contained in the optimal multicut by $\mathsf{OPT}(G,H)$. All of our discussion also applies to weighted network graphs.

2.1 Relaxations of MultiCut

In this section, we present a plethora of relaxations for MultiCut and discuss their properties and the relations between them. The proofs of theorems in this section can be found in appendix A (the organization of §A is essentially the same as the organization of this section).

LP relaxation. A function $d\colon X \times X \to \mathbb{R}_+$ is a *metric*[4] on X, if it satisfies the triangle inequality for all triples in X, i.e., $d(x,y) \leqslant d(x,z) + d(z,y)$ for any three points $x,y,z \in X$. Garg, Vazirani and Yannakakis [GVY96] consider the following linear programming relaxation of MultiCut.

$$\mathsf{LP}(G,H)\colon \quad \text{minimize} \quad \mathop{\mathbb{E}}_{(i,j)\in E(G)} d(i,j) \tag{3}$$

$$\text{subject to} \quad \forall_{(s,t)\in E(H)}\ d(s,t) = 1\,, \tag{4}$$

$$d \text{ is a metric on } V\,. \tag{5}$$

(We use $\mathsf{LP}(G,H)$ to refer to both the relaxation and its optimal objective value.) The integrality gap ratio of this relaxation as a function of the size of H is $\Theta(\log|E(H)|)$ [GVY96]. In section 2.3.2, we present (stronger) upper and lower bounds on the integrality gap in terms of different parameters of the demand graph H. In particular, we show that the gap remains $\Theta(\log|E(H)|)$ even on demand graphs with maximum degree 1.

SDP relaxation. Alternatively, we can consider the following natural SDP relaxation of MultiCut.

$$\mathsf{SDP}(G,H)\colon \quad \text{minimize} \quad \mathop{\mathbb{E}}_{(i,j)\in E(G)} \tfrac{1}{2}\|v_i - v_j\|^2 \tag{6}$$

$$\text{subject to} \quad \forall_{(s,t)\in E(H)}\ \langle v_s, v_t\rangle = 0\,, \tag{7}$$

$$\forall_{i\in V} \quad \|v_i\|^2 = 1\,. \tag{8}$$

This relaxation corresponds to the SDP relaxation of Max Cut used by Goemans and Williamson [GW95]. In contrast to the LP relaxation, this relaxation is sensitive to the maximum degree of the demand graph, denoted $\Delta := \Delta(H)$.

Theorem 2.1. *Let (G,H) be a MultiCut instance and $\varepsilon = \mathsf{SDP}(G,H)$. Then,* $\mathsf{OPT}(G,H) = O\left(\sqrt{\varepsilon \log(\Delta/\varepsilon)}\right)$.

[4] In this paper, we do not distinguish between semimetrics and proper metrics.

The proof of this theorem follows the analysis of the corresponding approximation algorithms for Unique Games [CMM06a, CMM06b]. The bound of Theorem 2.1 is almost tight (a consequence of known integrality gaps for Unique Games [KV05, KKMO04], together with the reduction from Unique Games to MultiCut).

Theorem 2.2. *For every $\varepsilon > 0$ and $\Delta = 2^{o(1/\varepsilon)}$, there exists* MultiCut *instances (G,H) with $\mathsf{SDP}(G,H) = \varepsilon$, $\Delta(H) = \Delta$, yet $\mathsf{OPT}(G,H) = \Omega(\sqrt{\varepsilon \log \Delta})$.*

Metric SDP relaxation. The approximation guarantees for the relaxations $\mathsf{LP}(G,H)$ and $\mathsf{SDP}(G,H)$ are incomparable. By combining the two relaxations, we can obtain an approximation ratio that is never worse than $O(\log n)$ but can be much better than $O(\log n)$ when demand-degree Δ is not too large.

$$
\begin{aligned}
\mathsf{SDP}_{\text{metric}}(G,H)\colon \quad & \text{minimize} && \mathop{\mathbb{E}}_{(i,j)\in E(G)} \tfrac{1}{2}\|v_i - v_j\|^2 && (9)\\
& \text{subject to} && \forall_{(s,t)\in E(H)}\ \langle v_s, v_t\rangle = 0\,, && (10)\\
& && \forall_{i\in V}\qquad \|v_i\|^2 = 1\,, && (11)\\
& && \qquad \|v_i - v_j\|^2 \text{ is a metric on } V\cup\{0\}\,. && (12)
\end{aligned}
$$

The constraint (12) requires that the distance function $d(i,j) = \|v_i - v_j\|^2$ satisfies the triangle inequality on the set $V \cup \{0\}$. Here, 0 is an additional point embedded in the origin, i.e., $v_0 = 0$. The following approximation guarantee follows by adapting the analysis of the Unique Games algorithm of Chlamtac, Makarychev and Makarychev [CMM06b].

Theorem 2.3. *Let (G,H) be a* MultiCut *instance on n vertices with $\varepsilon = \mathsf{SDP}_{\text{metric}}(G,H)$. Then, $\mathsf{OPT}(G,H) = O\left(\varepsilon\sqrt{\log n \log (\Delta/\varepsilon)}\right)$.*

Note that the approximation ratio in this theorem is never worse than $O(\log n)$ because we can assume $\Delta \leqslant n$ and $\varepsilon \geqslant 1/n^2$ (at least for unweighted graphs).

2.1.1 Intractable Relaxations

In the following, we introduce two very strong relaxations. We do know whether they can be solved efficiently. (Assuming the UGC no efficient algorithm can solve them.) We introduce these relaxations, because we find them helpful for understanding the previously introduced relaxations.

The first intractable relaxation has only a constant integrality gap. We use this relaxation to clarify the rounding problem of MultiCut (especially for the relaxation $\mathsf{SDP}(G,H)$ and $\mathsf{SDP}_{\text{metric}}(G,H)$).

The purpose of the second intractable relaxation is to estimate the behavior of $\mathsf{SDP}(G,H)$ when (G,H) is a product instance.

Characterization of MultiCut in terms of L_1 metrics. We say that $d\colon X\times X \to \mathbb{R}_+$ is an L_1 metric on a set X if there exists a multiplier α and a distribution over cuts $S \subseteq X$ such that $d(i,j) = \alpha \cdot \mathbb{P}\{i \text{ and } j \text{ are separated by } S\}$.

$$
\begin{aligned}
\mathsf{L}_1(G,H)\colon \quad & \text{minimize} && \mathop{\mathbb{E}}_{(i,j)\in E(G)} d(v_i, v_j) && (13)\\
& \text{subject to} && \forall_{(s,t)\in E(H)}\ d(v_s, v_t) = 1\,, && (14)\\
& && \forall_{i\in V}\qquad d(0,i) = 1/2\,, && (15)\\
& && \qquad d \text{ is an } L_1 \text{ metric on } V\cup\{0\}\,. && (16)
\end{aligned}
$$

This relaxation of MultiCut characterizes the $\mathsf{OPT}(G,H)$ up to a factor 2. (It is an interesting question whether this factor can be avoided.)

Theorem 2.4. *For every* MultiCut *instance (G,H), we have $\mathsf{OPT}(G,H) \leqslant 2\cdot \mathsf{L}_1(G,H)$.*

It is remarkable that constraint (15) plays a crucial role for this theorem. Without constraint (15) the integrality gap of the relaxation would be $\Omega(\log n)$ [ACMM05].

Non-negative SDP relaxation. We consider a strengthening of the relaxation $\mathsf{SDP}(G, H)$ where we add the constraint that all vectors lie in the non-negative orthant $\mathbb{R}^N_+$ of a sufficiently high dimensional Euclidean space (the dimension N is allowed to depend arbitrarily on the size of the instance (G, H)).

$$
\begin{aligned}
\mathsf{SDP}_+(G,H): \quad & \text{minimize} && \mathop{\mathbb{E}}_{(i,j)\in E(G)} \tfrac{1}{2}\|v_i - v_j\|^2 && (17)\\
& \text{subject to} && \forall_{(s,t)\in E(H)}\ \langle v_s, v_t\rangle = 0\,, && (18)\\
& && \forall_{i\in V} \qquad \|v_i\|^2 = 1\,, && (19)\\
& && \forall_{i\in V} \qquad v_i \in \mathbb{R}^N_+\,. && (20)
\end{aligned}
$$

The remarkable property of the this relaxation is that it has an approximation guarantee independent of the degree of the demand graph, and it differs from the relaxation $\mathsf{SDP}(G, H)$ by at most a factor of $O(\log (\Delta/\varepsilon))$.

Theorem 2.5. *For every* MultiCut *instance (G, H), we have* $\mathsf{OPT}(G, H) \leqslant O\left(\sqrt{\mathsf{SDP}_+(G, H)}\right)$.

Theorem 2.6. *For every* MultiCut *instance (G, H), we have* $\mathsf{SDP}_+(G, H) \leqslant \mathsf{SDP}(G, H) \cdot O(\log (\Delta/\varepsilon))$.

The first theorem follows, after a standard application of Cauchy–Schwarz, from (the proof of) Theorem 2.4. The second theorem is obtained by adapting a construction in [BHH+08]. Also notice that Theorem 2.1 follows immediately from these two theorems about $\mathsf{SDP}_+(G, H)$.

2.2 Reduction from Unique Games to MultiCut

(See appendix B for more details and proofs about this reduction.)

Γ Max 2 Lin variant of Unique Games. We consider a straight-forward reduction Φ (a generalization of the well-known reduction in [KRAR95]) that maps a Unique Games instance $\mathcal{U}$ to a multicut instance $(G, H) = \Phi(\mathcal{U})$. We analyze our reduction for a well-known special case of Unique Games, called Γ Max 2 Lin (sometimes also referred to as *linear unique games* or $\mathsf{E2LIN}(k)$). We need this additional assumption to ensure that the best multicut in $\Phi(\mathcal{U})$ is at least as good as the best labeling for the Unique Games instance $\mathcal{U}$.

We say a Unique Games instance $\mathcal{U}$ has Γ Max 2 Lin form, if the label set of $\mathcal{U}$ can be identified with the group $\mathbb{Z}_k$ in such a way that every permutation π_{uv} in $\mathcal{U}$ satisfies $\pi(i + s) = \pi_{uv}(i) + s \in \mathbb{Z}_k$ for all $s, i \in \mathbb{Z}_k$. In other words, π_{uv} encodes a constraint of the form $x_u - x_v = c_{uv} \in \mathbb{Z}_k$. Assuming this special structure of the permutations does not make the problem any easier [KKMO04].

Theorem 2.7. *Given a Γ Max 2 Lin instance $\mathcal{U}$, we can efficiently compute a multicut instance $(G, H) = \Phi(\mathcal{U})$ with $\varepsilon \leqslant \mathsf{OPT}(G, H) \leqslant 2\varepsilon$, where ε is the fraction of constraints violated by the optimal assignment for $\mathcal{U}$. The maximum degree of H is less than the alphabet size of $\mathcal{U}$.*

Combining this theorem and the $O(\log n)$-approximation for MultiCut [GVY96] yields a corresponding approximation for Γ Max 2 Lin. The same approximation is obtained by Gupta and Talwar [GT06] using different (more involved) techniques. (However, their LP relaxation for Unique Games is essentially the same as the LP relaxation for MultiCut used in [GVY96].) Their algorithm has the benefit that it works also for general Unique Games instances.

Corollary 2.8. *Given a Γ Max 2 Lin instance $\mathcal{U}$ such that an optimal assignment satisfies $1 - \varepsilon$ of the constraints, we can efficiently compute a labeling for $\mathcal{U}$ which satisfies at least $1 - O(\varepsilon \log n)$ of the constraints.*

2.3 Further Results and Discussions

2.3.1 Better SDP Approximation for Product Instances of MultiCut

In this section, we want to better approximation for product instances using the intractable relaxation $\mathsf{SDP}_+(G, H)$.

The following notion of product of MultiCut instances is motivated by the notion of parallel repetition of Unique Games instances.

Definition 2.9. The product of two multicut instance (G,H) and (G',H') on vertex sets V and V', respectively, is defined as $(G,H)\otimes(G',H') := (G\otimes_{\mathsf{AND}} G', H\otimes_{\mathsf{OR}} H')$. In the graph $G\otimes_{\mathsf{AND}} G'$, we connect (u,u') and (v,v') if both coordinates are either equal or adjacent in the original graphs. In the graph $H\otimes_{\mathsf{OR}} H'$, we connect (u,u') and (v,v') if one of the coordinates is adjacent in the original graph.

We can prove the following theorem for product instances.

Theorem 2.10. *Let (G,H) be a multicut instance with $\mathsf{SDP}(G,H)=\varepsilon$. Then, for $t\in\mathbb{N}$,*

$$\mathsf{OPT}\left(G^{\otimes_{AND} t}, H^{\otimes_{OR} t}\right) \leqslant O\left(\sqrt{t\cdot\varepsilon\log(\Delta/\varepsilon)}\right),$$

where Δ is the maximum degree of the demand graph H.

The proof of this theorem starts with a solution for the relaxation $\mathsf{SDP}(G,H)$. By Theorem 2.6, we can transform this solution to a solution for $\mathsf{SDP}_+(G,H)$ by increasing the cost by at most a factor $O(\log(\Delta/\varepsilon))$. By tensoring this solution for $\mathsf{SDP}_+(G,H)$, we obtain a solution for $\mathsf{SDP}_+(G^{\otimes_{\mathsf{AND}} t}, H^{\otimes_{\mathsf{OR}} t})$ of cost at most $t\cdot\varepsilon\log(\Delta/\varepsilon)$. Thus, by Theorem 2.5, there exists a multicut for the product instance $(G,H)^{\otimes t}$ of cost $O\left(\sqrt{t\cdot\varepsilon\log(\Delta/\varepsilon)}\right)$.

We claim that for "bipartite" MultiCut instances, Theorem 2.10 strongly improves the guarantee of Theorem 2.1 for suitable values of t. (We say that a MultiCut instance (G,H) is *bipartite* if there exists a bipartition (V_1,V_2) of V such that all edges of G go across the bipartition and all edges of H stay within their part of the bipartition.) The reason is the multiplicative property of the relaxation $\mathsf{SDP}(G,H)$ for bipartite instances.

Lemma 2.11 ([MS07, LM08]). *Let (G,H) be a bipartite multicut instance on the vertex set V. Then, for all $t\in\mathbb{N}$,*

$$\mathsf{SDP}\left(G^{\otimes_{AND} t}, H^{\otimes_{OR} t}\right) = 1-\left(1-\mathsf{SDP}(G,H)\right)^t \leqslant t\cdot\mathsf{SDP}(G,H).$$

This lemma shows that in Theorem 2.10, we have $t{\cdot}\varepsilon\approx\mathsf{SDP}(G^{\otimes_{\mathsf{AND}} t}, H^{\otimes_{\mathsf{OR}} t})$. Hence, if we would apply Theorem 2.1 on the product instance $(G,H)^{\otimes t}$, we would get the bound $\mathsf{OPT}(G,H)^{\otimes t}\leqslant O(t\sqrt{\varepsilon\log(\Delta/\varepsilon)})$, which is worse than the bound of Theorem 2.10 by a factor $\sqrt{t}$. (Here, we are using that $\Delta(H^{\otimes_{\mathsf{OR}} t})\geqslant\Delta(H)^t$.)

2.3.2 LP Relaxation and Parameters of the Demand Graph

(The proofs of the theorems in this section are in appendix A.4.)

A lower bound. First we note that restricting the degree of the demand graph does not improve the integrality gap of the LP relaxation.

Theorem 2.12. *For every $n\in\mathbb{N}$, there exists a multicut instance (G,H) on n vertices such that $\Delta(H)=1$, $\mathsf{LP}(G,H)=O(1/\log n)$, yet $\mathsf{OPT}(G,H)=\Omega(1)$.*

For this theorem, we can choose G as a random 2-lift of, say, a random 3-regular graph G_0 (see [AL02] for the notion of lifts of graphs). In the graph H, we match all vertices that correspond to the same vertex in G_0. A similar theorem is shown in [GT06].

An upper bound. Next we introduce a graph parameter $\alpha^*(H)$ that allows to prove a better bound on the integrality gap of LP relaxation. Let $\alpha^*(H) := \min_S \max_T |S\cap T|$, where S ranges over all vertex covers of H and T ranges over all independent sets of H. The integrality gap ratio of $\mathsf{LP}(G,H)$ is only logarithmic in $\alpha^*(H)$.

Theorem 2.13. *For every* MultiCut *instance (G,H), we have $\mathsf{OPT}(G,H)\leqslant\mathsf{LP}(G,H)\cdot O(\log\alpha^*(H))$.*

Notice that $\alpha^*(H)$ is less than both the vertex cover number of H and the independent set number of H. It is easy to see that $\alpha^*(H) < O(\log n)$ for most dense graphs H (say, H is drawn from the distribution $\mathcal{G}(n,1/2)$).

The proof of this theorem is by inspection of one of the known proofs [CKR01, GKL03] that $\mathsf{LP}(G,H)$ has integrality ratio $O(\log n)$.

2.3.3 Tightness of SDP Approximations

It is an interesting question whether in our approximation guarantees (Theorems 2.1, 2.3) the dependence on Δ/ε can be improved to just Δ. At least for the second guarantee it seems unlikely that an approximation algorithm based on the same relaxation can achieve such a guarantee. Suppose we could round an solution for $\mathsf{SDP}_{\text{metric}}(G, H)$ of value ε to a multicut of value $O(\varepsilon\sqrt{\log n \log \Delta})$. We claim that this bound would imply a $O(\sqrt{\log n})$-approximation even for arbitrary demand degrees. The reason is that one can reduce the degrees of the demand graph to 1 by splitting[5] every vertex in sufficiently many copies. To make sure that the resulting multicut instance is equivalent to the original one it suffices to connect all copies of a vertex by a clique in the demand graph. If we give enough weight to the edges in those cliques the optimal multicut will not separate any of the copies of a vertex.

On the other hand, the strong integrality gap for MultiCut [ACMM05] suggests that an algorithm based on $\mathsf{SDP}_{\text{metric}}$ cannot achieve an $O(\sqrt{\log n})$-approximation for general demand degrees.

3 Selected Proof

In this section we give very simple proof of an approximation guarantee for $\mathsf{SDP}(G, H)$ which is slightly weaker than the guarantee in Theorem 2.1.

Theorem 3.1. *For every* MultiCut *instance* (G, H), *if* $\mathsf{SDP}(G, H) = \varepsilon$ *then* $\mathsf{OPT}(G, H) \leqslant O\left(\sqrt{\varepsilon}\log(\Delta/\varepsilon)\right)$.

We prove this theorem by a natural rounding procedure. The analysis combines ideas of Goemans and Williamson [GW95] and Karger, Motwani and Sudan [KM94].

Proof. Let $v_1, \ldots, v_n \in \mathbb{R}^n$ be an optimal solution to $\mathsf{SDP}(G, H)$. Sample $r = \log(\Delta/\varepsilon)$ random hyperplanes through the origin. Let $S_1, \ldots, S_R \subseteq \mathbb{R}^n$ be the induced partition of $\mathbb{R}^n$ ($R = 2^r$ with probability 1). Let $\bar{M}$ be the set of pairs (i, j) such that v_i and v_j lie in different parts of the partition. For every edge $(i, j) \in E(G)$, we can upper bound the probability that the vectors of i and j lie in different parts by

$$\mathbb{P}\left\{(i, j) \in \bar{M}\right\} \leqslant r \cdot \|v_i - v_j\| \,.$$

For every demand pair $(s, t) \in E(H)$, we can upper bound the probability that the vectors of s and t lie in the same part by

$$\mathbb{P}\left\{(s, t) \notin \bar{M}\right\} = (1/2)^r \,.$$

Let V_0 be the set of vertices s such that the vector of a neighbor of s in H lies in the same part as v_s. We can upper bound the probability that a vertex s is in the set V_0 by

$$\mathbb{P}\left\{s \in V_0\right\} \leqslant \deg_H(s)(1/2)^r \,.$$

Let us consider the multicut that consists of the components $K_j = \{i \in V \setminus V_0 \mid v_i \in S_j\}$ and singleton components for each vertex in V_0. Let M be the set of edges in G that cross the multicut. We have

$$\mathbb{P}\left\{(i, j) \in M\right\} \leqslant \mathbb{P}\left\{(i, j) \in \bar{M}\right\} + \mathbb{P}\left\{i \in V_0\right\} + \mathbb{P}\left\{j \in V_0\right\} \leqslant r \cdot \|v_i - v_j\| + 2^{-r}(\deg_H(i) + \deg_H(j)) \,.$$

Hence, we can give the desired upper bound on $\mathsf{OPT}(G, H)$,

$$\tfrac{1}{|E(G)|}\,\mathbb{E}\,|M| \leqslant r \mathop{\mathbb{E}}_{(i,j)\in E(G)} \|v_i - v_j\| + 2\Delta \cdot 2^{-r} \leqslant r\sqrt{\varepsilon} + 2\Delta(H) \cdot 2^{-r} = \sqrt{\varepsilon}\log(\Delta/\varepsilon) + 2\Delta \cdot (\varepsilon/\Delta) = O(\sqrt{\varepsilon}\log(\Delta/\varepsilon)) \,.$$

□

[5] This idea of splitting vertices to reduce the degrees of the demand graph was communicated to us by Julia Chuzhoy after we presented our work to her.

References

[ACMM05] Amit Agarwal, Moses Charikar, Konstantin Makarychev, and Yury Makarychev. $O(\sqrt{\log n})$-approximation algorithms for min UnCut, min 2CNF deletion, and directed cut problems. In *STOC*, pages 573–581, 2005. 1, 2, 6, 9

[AKK+08] Sanjeev Arora, Subhash Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth K. Vishnoi. Unique games on expanding constraint graphs are easy: extended abstract. In *STOC*, pages 21–28, 2008. 3

[AL02] Alon Amit and Nathan Linial. Random graph coverings i: General theory and graph connectivity. *Combinatorica*, 22(1):1–18, 2002. 8, 14

[AR98] Yonatan Aumann and Yuval Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM J. Comput.*, 27(1):291–301, 1998. 1

[ARV04] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. In *STOC*, pages 222–231, 2004. 1, 12

[BHH+08] Boaz Barak, Moritz Hardt, Ishay Haviv, Anup Rao, Oded Regev, and David Steurer. Rounding parallel repetitions of unique games. In *FOCS*, pages 374–383, 2008. 2, 7, 14

[CKK+05] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. In *Proceedings of the Annual IEEE Conference on Computational Complexity*, number 20, pages 144–153, 2005. 1

[CKR01] Gruia Calinescu, Howard J. Karloff, and Yuval Rabani. Approximation algorithms for the 0-extension problem. In *SODA*, pages 8–16, 2001. 8, 15

[CMM06a] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for unique games. In *STOC*, pages 205–214, 2006. 1, 2, 4, 6

[CMM06b] Eden Chlamtac, Konstantin Makarychev, and Yury Makarychev. How to play unique games using embeddings. In *FOCS*, pages 687–696, 2006. 1, 2, 6, 12

[DS05] Irit Dinur and Shmuel Safra. On the hardness of approximating minimum vertex cover. *Annals of Mathematics*, 162(1):439–485, 2005. 1

[GKL03] Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS*, pages 534–543, 2003. 8, 15

[GT06] Anupam Gupta and Kunal Talwar. Approximating unique games. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 99–106, New York, NY, USA, 2006. ACM. 1, 7, 8

[GVY96] Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Approximate Max-Flow Min-(Multi)Cut theorems and their applications. *SIAM Journal on Computing*, 25(2):235–251, April 1996. 1, 5, 7

[GW95] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995. 5, 9

[Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the ACM Symposium on the Theory of Computing*, number 34, pages 767–775, 2002. 1

[KKMO04] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable CSPs? In *Annual Symposium on Foundations of Computer Science*, number 45, pages 146–154, 2004. 1, 6, 7

[KM94] David Karger and Rajeev Motwani. Approximate graph coloring by semidefinite programming. In *Journal of the ACM*, pages 2–13, 1994. 9

[KR03] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2-\varepsilon$. *Computational Complexity, Annual IEEE Conference on*, 0:379, 2003. 1

[KRAR95] Philip Klein, Satish Rao, Ajit Agrawal, and R. Ravi. An approximate max-flow min-cut relation for undirected multicommodity flow, with applications. *Combinatorica*, 15(2):187–202, June 1995. 7

[KV05] Subhash A. Khot and Nisheeth K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. In *FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 53–62, Washington, DC, USA, 2005. IEEE Computer Society. 1, 6

[Lee05] James R. Lee. On distance scales, embeddings, and efficient relaxations of the cut cone. In *SODA*, pages 92–101, 2005. 12

[LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. 1

[LM08] Troy Lee and Rajat Mittal. Product theorems via semidefinite programming. In *ICALP (1)*, pages 674–685, 2008. 8

[LR99] Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, 1999. 1

[MS07] Rajat Mittal and Mario Szegedy. Product rules in semidefinite programming. In *FCT*, pages 435–445, 2007. 8

[Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every csp? In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 245–254, New York, NY, USA, 2008. ACM. 1

[Raz08] Ran Raz. A counterexample to strong parallel repetition. *Foundations of Computer Science, Annual IEEE Symposium on*, 0:369–373, 2008. 2

[RS09] Prasad Raghavendra and David Steurer. Integrality gaps for strong SDP relaxations of unique games. In *FOCS*, 2009. To appear. 2

[Tre05] Luca Trevisan. Approximation algorithms for unique games. In *FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 197–205, Washington, DC, USA, 2005. IEEE Computer Society. 1

[Vaz01] Vijay V. Vazirani. *Approximation algorithms*. Springer-Verlag New York, Inc., New York, NY, USA, 2001. 1

A Relaxations of MultiCut

A.1 Preliminaries

The following lemma simplifies the analysis of many of the rounding algorithm for MultiCut. Essentially, this lemma is the core of the proof of Theorem 2.4.

Lemma A.1. *Let (G, H) be a* MultiCut *instance with vertex set V. Consider a distribution over independent sets S of H, i.e., $\forall (s,t) \in E(H)$. $\mathbb{P}\{s,t \in S\} = 0$. Then,*

$$\mathsf{OPT}(G,H) \leqslant \mathop{\mathbb{E}}_{(u,v)\in E(G)} \frac{\mathbb{P}\left\{(u,v) \in (S,\bar{S})\right\}}{\mathbb{P}\{u \in S \vee v \in S\}}\,.$$

Proof. Let $S_1, \ldots, S_r, \ldots$ be an infinite sequence of independent samples from this distribution. We extract a sequence of disjoint vertex sets $T_1, \ldots, T_r, \ldots$ as follows

$$T_r = S_r \setminus (S_{r-1} \cup \cdots \cup S_1)\,.$$

Since V is a finite set, almost all sets T_r are empty. If we discard the empty sets, we obtain a partition P of V. The constraint (A.1) implies that none of the sets T_r contain an edge of H. Hence, P separated all demand edges. It remains to bound the fraction of network edges separated by P. Let (u,v) be an edge in G. Let us condition on the event that S_r is the first set that contains either u or v. If $(u,v) \in M$, then it must be the case that $(u,v) \in (S_r, \bar{S}_r)$. Hence,

$$\mathbb{P}\{P(u) \neq P(v)\} \leqslant \frac{\mathbb{P}\left\{(u,v) \in (S,\bar{S})\right\}}{\mathbb{P}\{u \in S \text{ or } v \in S\}}\,.$$

We can conclude that

$$\mathsf{OPT}(G,H) \leqslant \mathop{\mathbb{E}}_{(u,v)\in E(G)} \mathbb{P}\{P(u) \neq P(v)\} < \mathop{\mathbb{E}}_{(u,v)\in E(G)} \frac{\mathbb{P}\left\{(u,v) \in (S,\bar{S})\right\}}{\mathbb{P}\{u \in S \text{ or } v \in S\}}\,.$$

$\square$

A.2 SDP Relaxations

A.2.1 Basic SDP Relaxation

Lemma A.2. *Let* (G, H) *be a multicut instance. Suppose* $\mathsf{SDP}(G, H) = \varepsilon$. *Then,*

$$\mathsf{OPT}(G, H) \leqslant O(\sqrt{\varepsilon \log(\Delta/\varepsilon)})\,.$$

Proof. Let $v_1, \ldots, v_n \in \mathbb{R}^n$ be an optimal solution for $\mathsf{SDP}(G, H)$ of value ε. Let g be a standard n-dimensional Gaussian vector, and let S be the set of vertices whose vectors have a projection on g larger than $\tau = C\sqrt{\log(\Delta/\varepsilon)}$,

$$S = \{i \in V \mid \langle g, v_i \rangle \geqslant \tau\}\ .$$

Let $N(\tau)$ denote the probability that a standard Gaussian variable is at least τ. The behavior of the Gaussian tail tells us $N(\tau) \ll \varepsilon/\Delta$. For every edge $(s, t) \in E(H)$, the projections $\langle g, v_s \rangle$ and $\langle g, v_t \rangle$ are independent and thus

$$\mathbb{P}\{s, t \in S\} = N(\tau)^2\,.$$

Let S_0 be the set of vertices that are in S and have a neighbor from H in S. Then, for every $s \in V$,

$$\mathbb{P}\{s \in S_0\} = \deg_H(s) N(\tau)^2\,.$$

Consider an edge $(i, j) \in E(G)$. Let $\varepsilon_{ij} = \frac{1}{2}\|v_i - v_j\|^2$. We can write $v_j = (1 - \varepsilon_{ij})v_i + \Theta(\sqrt{\varepsilon_{ij}})v_i^\perp$, where $v_i^\perp$ is a unit vector orthogonal to v_i. We have

$$\begin{aligned}
\mathbb{P}\{j \notin S \mid i \in S\} &= \mathbb{P}\left\{(1 - \varepsilon_{ij})\langle g, v_i \rangle + \Theta(\sqrt{\varepsilon_{ij}})\langle g, v_i^\perp \rangle < \tau \mid \langle g, v_i \rangle \geqslant \tau\right\} \\
&= O(\sqrt{\varepsilon_{ij}}\,\tau) \quad \text{(by [CMM06b, Lemma A.2])}\,.
\end{aligned}$$

And, therefore $\mathbb{P}\left\{(i, j) \in (S, \bar{S})\right\} \leqslant O(\sqrt{\varepsilon_{ij}})\tau N(\tau)$. Let $T = S \setminus S_0$. Then,

$$\begin{aligned}
\frac{\mathbb{P}\left\{(i, j) \in (T, \bar{T})\right\}}{\mathbb{P}\{i \in T \vee j \in T\}} &\leqslant \frac{\mathbb{P}\left\{(i, j) \in (S, \bar{S})\right\} + 2\Delta \cdot N(\tau)^2}{\mathbb{P}\{i \in S \vee j \in S\} - 2\Delta \cdot N(\tau)^2} \\
&\leqslant \frac{O(\sqrt{\varepsilon_{ij}})\tau N(\tau) + 2\Delta N(\tau)^2}{N(\tau) \cdot (1 - 2\Delta \cdot N(\tau))} \\
&= O(\sqrt{\varepsilon_{ij}})\tau + \varepsilon \qquad \text{(using } N(\tau) \ll \varepsilon/\Delta\text{)}
\end{aligned}$$

By Lemma A.1 there exists a multicut of size

$$\mathsf{OPT}(G, H) \leqslant \mathop{\mathbb{E}}_{(i,j)\in E(G)} O(\sqrt{\varepsilon_{ij} \log(\Delta/\varepsilon)} + \varepsilon) \leqslant O(\sqrt{\varepsilon \log(\Delta/\varepsilon)})$$

□

A.2.2 Metric SDP Relaxation

The following lemma is proved in [CMM06b] (based on [ARV04, Lee05]).

Lemma A.3. *Let* $V = [n]$ *and let* $\{v_1, \ldots, v_n\} \subseteq \mathbb{R}^n$ *be a collection of unit vectors satisfying the* ℓ_2^2*-triangle inequality (also including the origin). Then, for every* $m \geqslant 0$ *there exists a distribution over subsets* $S \subseteq V$ *such that*

– *for every vertex* $i \in V$,

$$\mathbb{P}\{i \in S\} = \alpha\,,$$

– *for any two vertices* $s, t \in V$ *with* $\langle v_s, v_t \rangle = 0$,

$$\mathbb{P}\{s, t \in S\} \leqslant \alpha \cdot \tfrac{1}{m},$$

– *for any two vertices* $i, j \in V$,

$$\mathbb{P}\left\{(i,j) \in (S, \bar{S})\right\} \leqslant \alpha \cdot O(\sqrt{\log m \log n}) \cdot \|v_i - v_j\|^2 .$$

Lemma A.4. *Let* (G, H) *be a multicut instance. Suppose* $\mathsf{SDP}_{\text{metric}}(G, H) = \varepsilon$*. Then,*

$$\mathsf{OPT}(G,H) \leqslant O\left(\varepsilon \cdot \sqrt{\log n} \cdot \sqrt{\log (\Delta/\varepsilon)}\right) .$$

Proof. Let $\{v_1, \ldots, v_n\} \subseteq \mathbb{R}^n$ be an optimal solution for $\mathsf{SDP}_{\text{metric}}(G, H)$ of value ε. Consider the distribution over subsets $S \subseteq V$ from the lemma above for $m = \Delta/\varepsilon$. We denote by S_0 the set of vertices in S that have a demand-neighbor in S. By the union bound, we have for every terminal $s \in V$,

$$\mathbb{P}\{s \in S_0\} \leqslant \Delta \cdot \alpha \cdot \tfrac{1}{m} = \alpha \cdot \varepsilon .$$

Let T be the set-valued random variable defined by $S \setminus S_0$. For every network edge (i, j), we have

$$\frac{\mathbb{P}\left\{(i,j) \in (T, \bar{T})\right\}}{\mathbb{P}\{i \in T \ \vee \ j \in T\}} \leqslant \frac{\mathbb{P}\left\{(i,j) \in (S, \bar{S})\right\} + \mathbb{P}\{i \in S_0\} + \mathbb{P}\{j \in S_0\}}{\alpha} \leqslant \frac{\alpha \cdot O(\sqrt{\log m \log n}) \cdot \|v_i - v_j\|^2 + 2\alpha\varepsilon}{\alpha}$$
$$\leqslant O(\varepsilon + \sqrt{\log m \log n}\|v_i - v_j\|^2) . \quad (21)$$

Using Lemma A.1, we can finish the proof

$$\mathsf{OPT}(G,H) \leqslant \mathop{\mathbb{E}}_{(i,j)\in E(G)} \frac{\mathbb{P}\left\{(i,j) \in (T, \bar{T})\right\}}{\mathbb{P}\{i \in T \ \vee \ j \in T\}} \leqslant \mathop{\mathbb{E}}_{(i,j)\in E(G)} O(\varepsilon + \sqrt{\log m \log n}\|v_i - v_j\|^2) = O(\varepsilon\sqrt{\log m \log n}) .$$

□

A.3 Intractable Relaxations

A.3.1 Characterization of MultiCut by L_1 metrics

Proof. [Proof of Theorem 2.4] Let d be an optimal solution to $\mathsf{L}_1(G, H)$. We can express d as a distribution over subsets $S \subseteq V(G)$. It is easy to see that constraint (15) implies that the sets in the support of this distribution are independent sets of H. From Lemma A.1, we get the desired bound on the multicut value $\mathsf{OPT}(G, H)$,

$$\begin{aligned}
\mathsf{OPT}(G,H) &\leqslant \mathop{\mathbb{E}}_{(u,v)\in E(G)} \frac{\mathbb{P}\left\{(u,v) \in (S, \bar{S})\right\}}{\mathbb{P}\{u \in S \vee v \in S\}} \\
&\leqslant \mathop{\mathbb{E}}_{(u,v)\in E(G)} \frac{\mathbb{P}\left\{(u,v) \in (S, \bar{S})\right\}}{\max\{\mathbb{P}\{u \in S\}, \mathbb{P}\{v \in S\}\}} \\
&= \mathop{\mathbb{E}}_{(u,v)\in E(G)} \frac{\alpha d(u,v)}{\alpha/2} = 2 \cdot \psi_1(G,H) .
\end{aligned}$$

□

A.3.2 Non-negative SDP Relaxation

Theorem 2.6 follows from the following lemma, which is implicit in [BHH+08].

Lemma A.5. *[BHH+08] Let (G, H) be a multicut instance and let $v_1, \ldots, v_n \in \mathbb{R}^n$ be an optimal solution to $\mathsf{SDP}(G, H)$ of value ε. Then, there exists vectors $v'_1, \ldots, v'_n$ with only non-negative coordinates such that*

1. *for all $i \in [n]$, we have $\|v'_i\|^2 = 1$,*
2. *for all $(s, t) \in E(H)$, we have $\langle v'_s, v'_t\rangle = 0$.*
3. *for all $(i, j) \in E(G)$, we have $\frac{1}{2}\|v'_i - v'_j\|^2 \leqslant O(\log(\Delta/\varepsilon))\|v'_i - v'_j\|^2 + O(\varepsilon)$.*

Theorem 2.5 follows from the following lemma.

Lemma A.6. *Suppose $v_1, \ldots, v_n \in \mathbb{R}^d_+$ is a solution for $\mathsf{SDP}_+(G, H)$ of cost ε. Then,*

$$\mathsf{OPT}(G, H) \leqslant 2\sqrt{2\varepsilon}\,.$$

Proof. Let v_i^2 denote the vectors obtained from v_i by squaring every coordinate. We consider the L_1-metric d on $\{0, 1, \ldots, n\}$ defined by $d(i, j) = \frac{1}{2}\|v_i^2 - v_j^2\|_1$ and $d(0, i) = \frac{1}{2}\|v_i^2\|_1 = 1/2$. We have

$$2d(i, j) = \|v_i^2 - v_j^2\|_1 = \sum_{r=1}^{d} \left|v_i(r) - v_j(r)\right| \left|v_i(r) + v_j(r)\right| \leqslant \|v_i - v_j\|_2 \cdot \|v_i + v_j\|_2 \leqslant 2\|v_i - v_j\|_2\,.$$

For every demand pair $(s, t) \in E(H)$, the condition $\langle v_s, v_t\rangle = 0$ shows that the support of the non-negative vectors v_s and v_t is disjoint. Therefore, $d(s, t) = \frac{1}{2}\|v_s^2 - v_t^2\|_1 = \frac{1}{2}\|v_s^2\|_1 + \frac{1}{2}\|v_t^2\|_1 = 1$. It follows that d is a solution for $\mathsf{L}_1(G, H)$. Hence, by Theorem 2.4,

$$\mathsf{OPT}(G, H) \leqslant 2 \mathop{\mathbb{E}}_{(i,j)\in E(G)} d(i, j) \leqslant 2 \mathop{\mathbb{E}}_{(i,j)\in E(G)} \|v_i - v_j\|_2 \leqslant 2\left(2 \mathop{\mathbb{E}}_{(i,j)\in E(G)} \tfrac{1}{2}\|v_i - v_j\|_2^2\right)^{1/2} = 2\sqrt{2\varepsilon}\,.$$

□

A.4 LP Relaxation

A.4.1 A lower bound

Here we present a proof of Theorem 2.12.

Proof. Let G_0 be regular graph with girth $\Omega(\log n)$ and max-cut value at most $2/3$ (a random 3-regular graph works well). We consider the unique game on G_0 that corresponds to the max-cut problem, that is, for every edge (u, v) of G_0, we have the permutation $\pi_{uv}\colon \{0, 1\} \to \{0, 1\}$ with $\pi_{uv}(b) = 1 - b$. It is easy to see that the value of this unique game is equal to the fraction of edges contained in the max-cut of G_0. Hence, by the choice of G_0, the value of this unique game is at most $2/3$. Let (G, H) be the multicut instance obtained by applying the reduction in §B to the unique game on G_0. Note that H is just a matching because the unique game has an alphabet of size 2. By Lemma B.2, we have $1 - 2\mathsf{OPT}(G, H) \leqslant \mathsf{VAL}(G_0) \leqslant 2/3$. Thus $\mathsf{OPT}(G, H) \geqslant 1/6$.

We claim that $\mathsf{LP}(G, H) = O(1/\log n)$. Note that this claim together with the fact $\mathsf{OPT}(G, H) \geqslant 1/6$ implies the lemma. Let d be the shortest path metric of G scaled by a factor $1/\mathrm{girth}(G_0) = O(1/\log n)$. Since every edge of G has length $O(1/\log n)$ in this metric, it follows that the objective value for $\mathsf{LP}(G, H)$ achieved by d is $O(1/\log n)$. It remains to show that d satisfies all constraints of $\mathsf{LP}(G, H)$. Let (s, t) be any demand pair. Let P be the shortest path from s to t in G. Note that $d(s, t) = \mathrm{length}(P)/\mathrm{girth}(G_0)$. By the construction[6] of G, the path P corresponds to cycle in G_0 of the same length. Thus, $\mathrm{length}(P) \geqslant \mathrm{girth}(G_0)$ and $d(s, t) \geqslant 1$.

□

[6] In graph theory, this construction is known as "lifting" [AL02]. In this language, G would be called a 2-lift of G_0. There is a 2-to-1 correspondence between the paths of G and G_0.

A.4.2 An upper bound

In this section we prove Theorem 2.13.

Proof. Let the metric d on $V := V(G)$ be an optimal solution to $\mathsf{LP}(G, H)$. Suppose $\alpha^*(H) = \max_T |V^* \cap T|$ for a vertex cover V^* of H. (Recall that T ranges over all independent sets of H.) We may assume $V^* = [n]$. We construct a multicut as follows (the construction is from [CKR01], the analysis is along the lines of [GKL03]):

1. Pick a radius r uniformly at random from the interval $\left[\frac{1}{6}, \frac{2}{6}\right]$.

2. Pick a random ordering $\pi\colon V^* \to [n]$ of the non-isolated vertices of H.

3. For every vertex s with $\deg_H(s) \geqslant 1$, define

$$S_s := \{u \in V \mid d(s,u) \leqslant r \text{ and } d(t,u) > r \text{ for all } t \in V^* \text{ with } \pi(t) < \pi(s)\}$$

4. Output the partition P induced by the sets $S_1, \ldots, S_n$.

Since every cluster of the partition P has diameter at most $2/6$, every demand edge is separated by P.
For every vertex $u \in V$, let B_u denote the set of vertices $s \in V^*$ such that $\frac{1}{6} \leqslant d(s,u) \leqslant \frac{2}{6}$.
We claim that for any two vertices $u, v \in V$,

$$\mathbb{P}\{P(u) \neq P(v)\} \leqslant O\left(\log |B_u \cup B_v|\right) \cdot d(u,v) \tag{22}$$

Let us first show that the claim implies the theorem. Notice that any two vertices in B_u are at distance at most $\frac{4}{6} < 1$. Hence, $B_u \subseteq V^*$ does not contain a demand edge. Thus, $|B_u| \leqslant \alpha^*(H)$.

We can now estimate the fraction of edges in the multicut given by P as follows:

$$\mathop{\mathbb{E}}_{(u,v)\in E(G)} \mathbb{P}\{P(u) \neq P(v)\} \overset{(22)}{\leqslant} \mathop{\mathbb{E}}_{(u,v)\in E(G)} O\left(\log |B_u \cup B_v|\right) \cdot d(u,v) = O\left(\log |B_u \cup B_v|\right) \cdot \mathsf{LP}(G,H)\,.$$

It remains to prove the claim (22). This claim follows from the proof of [GKL03, Theorem 3.2]. We omit the details from this preliminary version of the paper. ◻

B Reduction from Unique Games to MultiCut

The reason why Γ Max 2 Lin instances are more convenient to start with is due to the following structure. Given a labeling Λ one can define a labeling $\Lambda + i$ for every $i \in \mathbb{Z}_k$ as follows: $(\Lambda + i)(v) := \Lambda(v) + i$. Then, $\mathsf{VAL}(\Lambda) = \mathsf{VAL}(\Lambda + i)$ for every $i \in \mathbb{Z}_k$. We will often abuse notation and use $[k]$ in place of $\mathbb{Z}_k$.

The Reduction. Let $\mathcal{U} = (G(V,E), [k], \{\pi_{uv}\}_{(u,v)\in E})$ be an instance of Γ Max 2 Lin. Consider the following instance of MultiCut derived from it. The vertex set of the MultiCut instance is $V' := V \times [k]$ and vertices are labeled by (v, i), where v is a vertex in the Γ Max 2 Lin instance and i is a potential label to v. The network graph G' consists of edges between (v, i) and (w, j) iff (v, w) is an edge in the Γ Max 2 Lin instance and $\pi_{vw}(i) = j$. The demand graph H' consists of edges (v, i) and (v, j) for every $v \in V$ and every $i \neq j \in [k]$. It follows from the reduction that the demand graph consists of n cliques of size k. Hence, the maximum degree of the demand graph is k.

The proof of Theorem 2.7 follows from the following two lemmata. The first lemma is straightforward though it uses crucially the fact that the instance $\mathcal{U}$ has Γ Max 2 Lin form.

Lemma B.1. *Suppose there is a labeling Λ of $\mathcal{U}$ such that $\mathsf{VAL}(\Lambda) \geqslant 1 - \varepsilon$, then there is multicut in (G', H') of value at most ε.*

Proof. We partition the vertex set V' into k parts using the labeling Λ of V: Let $S_c \subseteq V'$ be the set of vertices (v, i) with $i = \Lambda(v) + c$. The collection $P := \{S_1, \ldots, S_k\}$ forms a partition of V'. Let us compute the fraction of edges that are cut by the partition P,

$$\begin{aligned}
\mathop{\mathbb{P}}_{(x,y)\in E(G')} \{P(x) \neq P(y)\} &= \mathop{\mathbb{E}}_{i\in\mathbb{Z}_k} \mathop{\mathbb{P}}_{(u,v)\in E(G)} \Big\{P\big((u,i)\big) \neq P\big((v, \pi_{uv}(i))\big)\Big\} \quad \text{(by construction of } G'\text{)}\\
&= \mathop{\mathbb{E}}_{i\in\mathbb{Z}_k} \mathop{\mathbb{P}}_{(u,v)\in E(G)} \{\Lambda(u) - i \neq \Lambda(v) - \pi_{uv}(i)\} \quad \text{(by definition of } P\text{)}\\
&= \mathop{\mathbb{P}}_{(u,v)\in E(G)} \{\Lambda(u) \neq \pi_{uv}(\Lambda(v))\} \quad \text{(using } \pi_{uv}(i) = \pi_{uv}(\Lambda(v)) - \Lambda(v) + i\text{)}\\
&= \mathsf{VAL}(\Lambda)\,.
\end{aligned}$$

□

The second lemma is the more interesting direction of the Theorem 2.7. Here, starting from a multicut in (G', H') of value at most ε, which separates all demand pairs, we construct a labeling for $\mathcal{U}$ where at most ε fraction of constraints are not satisfied.

Lemma B.2. *Suppose that there is multicut in (G', H') of value at most ε, then there is a labeling for $\mathcal{U}$ of value at least $1 - 2\varepsilon$.*

Proof. We present a probabilistic construction similar to the proof of Lemma A.1 Let P be a partition of V' that separates all edges of H' and only an ε fraction of the edges of G'. We consider the distribution over subsets $S \subseteq V'$ obtained by choosing uniformly at random a cluster of P. Let $S_1, \ldots, S_r, \ldots$ be infinite sequence of independent random subsets of V', each chosen from this distribution.

We extract a labeling $\Lambda\colon V \to [k]$ from this sequence of sets in the following way: For a vertex $u \in V$, let $S_{r(u)}$ be the first set in the sequence that contains a vertex (u, i) for some label $i \in [k]$. Since P is a multicut, the label $i \in [k]$ such that (u, i) is in the set $S_{r(u)}$ is unique. We assign this label to vertex u.

Let us estimate the fraction of constraints violated by this labeling. Let (u, v) be an edge in the unique game $\mathcal{U}$. For $r \in \mathbb{N}$, let us condition on the event that S_r is the first set that contains a vertex $(u, \cdot)$ or $(v, \cdot)$. Now, if Λ violates (u, v), then it must be the case that one of the network edges $(u, i) \sim (v, \pi_{uv}(i))$ in G' is cut by the set S_r. Hence,

$$\mathbb{P}\{\Lambda(v) \neq \pi_{uv}(\Lambda(u))\} \leqslant \frac{\mathbb{P}\Big\{\exists i.\ \big((u,i), (v, \pi_{uv}(i))\big) \in (S, \bar{S})\Big\}}{\mathbb{P}\{\exists i.\ (u,i) \in S \text{ or } \exists i.\ (v,i) \in S\}} \leqslant \frac{\sum_{i=1}^{k} \mathbb{P}\Big\{\big((u,i), (v, \pi_{uv}(i))\big) \in (S, \bar{S})\Big\}}{\max\Big\{\mathbb{P}\{\exists i.\ (u,i) \in S\}, \mathbb{P}\{\exists i.\ (v,i) \in S\}\Big\}}\,. \tag{23}$$

Suppose that the multicut P separates a ε_{uv} fraction of the network edges of the form $(u, \cdot)$ and $(v, \cdot)$ vertices. Then, the right-hand side of (23) equals $2\varepsilon_{uv}$ (the neumerator is $k\varepsilon_{uv} \cdot 2/|P|$ and the denominator is $k/|P|$, where $|P|$ is the number of clusters in the partition P). We can conclude that

$$\mathsf{VAL}(\Lambda) \geqslant 1 - \mathop{\mathbb{E}}_{(u,v)\in E(G)} \mathbb{P}\{\Lambda(v) \neq \pi_{uv}(\Lambda(u))\} \geqslant 1 - \mathop{\mathbb{E}}_{(u,v)\in E(G)} 2\varepsilon_{uv} = 1 - 2\varepsilon\,.$$

□

ECCC ISSN 1433-8092

http://eccc.hpi-web.de