

On the Unique Games Conjecture

Subhash Khot
Courant Institute of Mathematical Sciences, NYU *

Abstract

This article surveys recently discovered connections between the Unique Games Conjecture and computational complexity, algorithms, discrete Fourier analysis, and geometry.

1 Introduction

Since it was proposed in 2002 by the author [60], the Unique Games Conjecture (UGC) has found many connections between computational complexity, algorithms, analysis, and geometry. This article aims to give a survey of these connections. We first briefly describe how these connections play out, which are then further explored in the subsequent sections.

Inapproximability: The main motivation for the conjecture is to prove inapproximability results for NP-complete problems that researchers have been unable to prove otherwise. The UGC states that a specific computational problem called the **Unique Game** is inapproximable. A gap-preserving reduction from the **Unique Game** problem then implies inapproximability results for other NP-complete problems. Such a reduction was exhibited in [60] for the **Min-2SAT-Deletion** problem and it was soon followed by a flurry of reductions for various problems, in some cases using variants of the conjecture.

Discrete Fourier Analysis: The inapproximability reductions from the **Unique Game** problem often use gadgets constructed from a boolean hypercube (also used with a great success in earlier works [18, 49, 51]). The reductions can be alternately viewed as constructions of Probabilistically Checkable Proofs (PCPs) and the gadgets can be viewed as probabilistic checking procedures to check whether a given codeword is a *Long Code*. Fourier analytic theorems on hypercube play a crucial role in ensuring that the gadgets indeed “work”. Applications to inapproximability actually lead to some new Fourier analytic theorems.

Geometry: The task of proving a Fourier analytic theorem can sometimes be reduced to an isoperimetry type question in geometry. This step relies on powerful invariance style theorems [98, 83, 24, 82], some of which were motivated in a large part by their application to inapproximability. The geometric questions, in turn, are either already studied or are new, and many of these remain challenging open questions.

Integrality Gaps: For many problems, the UGC rules out *every* polynomial time algorithm to compute a good approximate solution. One can investigate a less ambitious question: can we rule out algorithms based on a specific linear or semi-definite programming relaxation? This amounts to the so-called *integrality gap* constructions that are explicit combinatorial constructions, often with geometric flavor. It was demonstrated in [71] and subsequent papers that the reduction from

*Supported by NSF CAREER grant CCF-0833228, NSF Expeditions grant CCF-0832795, and BSF grant 2008059.

the Unique Game problem to a target problem can in fact be used to construct an (unconditional, explicit) integrality gap instance for the target problem. This strategy was used in [71] to resolve certain questions in metric geometry. A result of Raghavendra [90] shows duality between approximation algorithms and inapproximability reductions in the context of constraint satisfaction problems: a natural semi-definite programming relaxation leads to an algorithm and an integrality gap instance to the same relaxation leads to an inapproximability reduction.

Algorithms and Parallel Repetition: Attempts to prove or disprove the UGC have led to some very nice algorithms, a connection to the small set expansion problem in graphs, a deeper understanding of the parallel repetition theorem, and solution to a tiling problem in Euclidean space. Some of these works demonstrate that if the conjecture were true, there must be a certain trade-off between quantitative parameters involved, that there can be no proof along a certain strategy, and that if pushed to certain extremes, the conjecture becomes false.

Many of the aforementioned developments were quite unexpected (to the author at least). One hopes that the future progress culminates in a proof or a refutation of the conjecture, and irrespective of the outcome, new techniques and unconditional results keep getting discovered. The reader is referred to [104, 61, 62, 85] for surveys on related topics.

2 Preliminary Background

2.1 Approximation Algorithms and Inapproximability

Let \mathcal{I} denote an NP-complete problem. For an instance I of the problem with input size N ,¹ let $\text{OPT}(I)$ denote the value of the optimal solution. For a specific polynomial time approximation algorithm, let $\text{ALG}(I)$ denote the value of the solution that the algorithm finds (or its expected value if the algorithm is randomized). Let $C > 1$ be a parameter that could be a function of N .

Definition 2.1 *An algorithm achieves an approximation factor of C if on every instance I ,*

$$\begin{aligned} \text{ALG}(I) &\geq \text{OPT}(I)/C && \text{if } \mathcal{I} \text{ is a maximization problem,} \\ \text{ALG}(I) &\leq C \cdot \text{OPT}(I) && \text{if } \mathcal{I} \text{ is a minimization problem.} \end{aligned}$$

A maximization problem \mathcal{I} is proved to be inapproximable by giving a reduction from a canonical NP-complete problem such as 3SAT to a *gap version* of \mathcal{I} . A (c, s) -gap version of the problem, denoted as $\text{Gap}\mathcal{I}_{c,s}$ is a promise problem where either $\text{OPT}(I) \geq c$ or $\text{OPT}(I) \leq s$. Suppose there is a polynomial time reduction from 3SAT to $\text{Gap}\mathcal{I}_{c,s}$ for some $0 < s < c$, i.e. a reduction that maps a 3SAT formula ϕ to an instance I of the problem \mathcal{I} such that:

- (YES Case): If ϕ has a satisfying assignment, then $\text{OPT}(I) \geq c$.
- (NO Case): If ϕ has no satisfying assignment, then $\text{OPT}(I) \leq s$.

Such a reduction implies that if there were an algorithm with approximation factor strictly less than c/s for problem \mathcal{I} , then it would enable one to efficiently decide whether a 3SAT formula is satisfiable, and hence $\text{P} = \text{NP}$. Inapproximability results for minimization problems can be proved in a similar way.

¹In this article, N denotes the size of the input, and n is reserved to denote the dimension of a boolean hypercube as well as the number of “labels” for the Unique Game problem.

2.2 The PCP Theorem

In practice, a reduction as described above is actually a sequence of (potentially very involved) reductions. The first reduction in the sequence is the famous *PCP Theorem* [37, 11, 9] which can be phrased as a reduction from 3SAT to a gap version of 3SAT. For a 3SAT formula ϕ , let $\text{OPT}(\phi)$ denote the maximum fraction of clauses that can be satisfied by any assignment. Thus $\text{OPT}(\phi) = 1$ if and only if ϕ is satisfiable. The PCP Theorem states that there is a universal constant $\alpha^* < 1$ and a polynomial time reduction that maps a 3SAT instance ϕ to another 3SAT instance ψ such that:

- (YES Case): If $\text{OPT}(\phi) = 1$, then $\text{OPT}(\psi) = 1$.
- (NO Case): If $\text{OPT}(\phi) < 1$, then $\text{OPT}(\psi) \leq \alpha^*$.

We stated the PCP Theorem as a combinatorial reduction. There is an equivalent formulation of it in terms of *proof checking*. The theorem states that every NP statement has a polynomial size proof that can be checked by a probabilistic polynomial time verifier by reading only a constant number of bits in the proof! The verifier has the completeness (i.e. the YES Case) and the soundness (i.e. the NO Case) property: every correct statement has a proof that is accepted with probability 1 and every proof of an incorrect statement is accepted with only a small probability, say at most 1%. The equivalence between the two views, namely reduction versus proof checking, is simple but illuminating, and has influenced much of the work in this area.

2.3 Towards Optimal Inapproximability Results

The PCP Theorem shows that Max-3SAT is inapproximable within factor $\frac{1}{\alpha^*} > 1$. By results of Papadimitriou and Yannakakis [89], the same holds for every problem in the class MAX-SNP. After the discovery of the PCP Theorem, the focus shifted to proving optimal results, that is proving approximability and inapproximability results that match with each other. In the author’s opinion, the most influential developments in this direction were the introduction of the Label Cover (a.k.a. 2-Prover-1-Round Game) problem [4], Raz’s Parallel Repetition Theorem [96], the introduction of the Long Code and the basic reduction framework of Long Code based PCPs [18], and Håstad’s use of Fourier analysis in analyzing the Long Code [49, 51]. Specifically, Håstad’s work gives optimal inapproximability results for 3SAT and the Clique problem. He shows that for every $\varepsilon > 0$, $\text{Gap3SAT}_{1, \frac{7}{8} + \varepsilon}$ is NP-hard and also that on N -vertex graphs, $\text{GapClique}_{N^{1-\varepsilon}, N^\varepsilon}$ is NP-hard. In words, given a satisfiable 3SAT formula, it is hard to find an assignment that satisfies $\frac{7}{8} + \varepsilon$ fraction of the clauses (a random assignment is $\frac{7}{8}$ -satisfying in expectation). Also, given an N -vertex graph that has a clique of size $N^{1-\varepsilon}$, it is hard to find a clique of size N^ε .

We give a quick overview of these developments which sets up the motivation for formulating the Unique Games Conjecture. We start with the (somewhat cumbersome) definition of the 2-Prover-1-Round Game problem.²

Definition 2.2 A 2-Prover-1-Round Game $\mathcal{U}_{2p1r}(G(V, W, E), [m], [n], \{\pi_e | e \in E\})$ is a constraint satisfaction problem defined as follows: $G(V, W, E)$ is a bipartite graph whose vertices represent variables and edges represent the constraints. The goal is to assign to each vertex in V a label from the set $[m]$ and to each vertex in W a label from the set $[n]$. The constraint on an edge $e = (v, w) \in E$ is described by a “projection” $\pi_e : [m] \mapsto [n]$. The projection is onto, may depend on

²The 2-Prover-1-Round games allow more general constraints than projections. The projection property is crucial for inapproximability applications and we restrict to this special case throughout this article.

the edge, and $m \geq n$. A labeling $L : V \mapsto [m], L : W \mapsto [n]$ satisfies the constraint on edge $e = (v, w)$ if and only if $\pi_e(L(v)) = L(w)$. Let $\text{OPT}(\mathcal{U}_{2p1r})$ denote the maximum fraction of constraints that can be satisfied by any labeling:

$$\text{OPT}(\mathcal{U}_{2p1r}) := \max_{L:V \mapsto [m], L:W \mapsto [n]} \frac{1}{|E|} \cdot |\{e \in E \mid L \text{ satisfies } e\}|.$$

The 2-Prover-1-Round Game problem is to find an (approximately) optimal labeling to a given instance. The term *game* derives from an equivalent formulation in terms of 2-Prover-1-Round games. Given an instance $\mathcal{U}_{2p1r}(G(V, W, E), [m], [n], \{\pi_e \mid e \in E\})$, consider the following game: there are two provers P_1 and P_2 and a probabilistic polynomial time verifier. The verifier picks an edge $e = (v, w)$ at random, sends the vertex v to prover P_1 and the vertex w to prover P_2 . The provers are supposed to answer with a label in $[m]$ and $[n]$ respectively. The verifier accepts if and only if $\pi_e(i) = j$ where i and j are the answers of the two provers respectively. The strategies of the provers correspond to a labeling $L : V \mapsto [m], L : W \mapsto [n]$. The *value* of the game is the maximum over all prover strategies, the probability that the verifier accepts and it is clearly the same as $\text{OPT}(\mathcal{U}_{2p1r})$. Thus the constraint satisfaction view and the 2-Prover-1-Round view are equivalent.

We are interested in the case where the sizes m and n of label sets are constant and the size of the constraint graph is the growing input size. The PCP Theorem implies that the gap version of the problem is NP-hard and then the gap can be amplified using parallel repetition theorem.

Theorem 2.3 (PCP Theorem + Raz’s Parallel Repetition Theorem:) *For every $\delta > 0$, $\text{Gap2P1R}_{1,\delta}$ is NP-hard on instances with label sets of size $\text{poly}(1/\delta)$. Specifically, there exists an absolute constant C such that for every $\delta > 0$, $\mathcal{U}_{2p1r}(G(V, W, E), [m], [n], \{\pi_e \mid e \in E\})$, $m = (1/\delta)^C$, it is NP-hard to distinguish between:*

- *YES Case:* $\text{OPT}(\mathcal{U}_{2p1r}) = 1$.
- *NO Case:* $\text{OPT}(\mathcal{U}_{2p1r}) \leq \delta$.

Many of the inapproximability results (e.g. [18, 49, 51, 47, 100]) are obtained by a reduction from the 2-Prover-1-Round Game in Theorem 2.3. Suppose one desires a reduction from $\text{Gap2P1R}_{1,\delta}$ to $\text{Gap}\mathcal{I}_{c,s}$ for a maximization problem \mathcal{I} . The core of such a reduction involves a gadget built from a boolean hypercube and the fact that the gadget “works” relies on Fourier analysis of boolean functions on the hypercube. One can also view the reduction as a PCP construction and the gadget as an encoding scheme along with a testing procedure to check whether a given string is indeed a codeword. An extremely useful encoding scheme is the so-called Long Code which is same as a *dictatorship* function on the hypercube. A dictatorship function $f : \{-1, 1\}^n \mapsto \{-1, 1\}$ is a function that depends only on one co-ordinate, i.e. $f(x) = x_{i_0}$ for some fixed $i_0 \in [n]$. The truth-table of this function is thought of as an encoding of index $i_0 \in [n]$ and this precisely is the Long Code. One desires a testing procedure such that a dictatorship function passes the test with probability at least c whereas any function that is *far from being a dictatorship* passes the test with probability at most s . The gap c/s essentially translates to the gap in the $\text{Gap}\mathcal{I}_{c,s}$ instance constructed by the reduction. Along with the quantitative gap, the nature of the test (number of queries, the acceptance predicate etc.) is also important and is dictated by the specific target problem \mathcal{I} .

The PCP, at a high level, involves replacing every vertex of the 2-Prover-1-Round Game by a copy of a boolean hypercube (i.e. the gadget). If $\mathcal{U}_{2p1r}(G(V, W, E), [m], [n], \{\pi_e \mid e \in E\})$ denotes the 2-Prover-1-Round Game instance, then every $v \in V$ is replaced by an m -dimensional hypercube and every $w \in W$ is replaced by an n -dimensional hypercube. Thus the labels are identified with

the hypercube dimensions. The PCP proof consists of truth-tables of boolean functions on these hypercubes. An honest proof is supposed to write down truth-tables of dictatorship functions, where for each hypercube, the co-ordinate relevant to its dictatorship function is the supposed label of the 2-Prover-1-Round Game vertex. The overall PCP testing has two parts: *codeword testing* that checks that each boolean function, by itself, is (close to) a dictatorship and *consistency testing* that checks that for $e = (v, w) \in E$, the boolean functions on hypercubes corresponding to v and w are dictatorships of some co-ordinates $i \in [m]$ and $j \in [n]$ respectively such that $\pi_e(i) = j$. Often, these two parts are combined together into a single test.

2.4 The Unique Games Conjecture

The overview of a typical PCP described above is a little too simplistic. There are often serious technical challenges to overcome and even though the above PCP strategy succeeds in some applications (e.g. Max-3SAT, Clique, Hypergraph Coloring [49, 51, 47, 100]), researchers have been unable to make it work for many other applications, e.g. Vertex Cover, MaxCut, Min-2SAT-Deletion, and Graph Coloring. An immediate observation is that in the first set of (successful) applications, the PCPs are allowed to make three or more queries, whereas in the second set of (unsuccessful) applications, the PCPs are allowed to make only two queries and two-query PCPs are potentially very weak. The nature of the PCP test (including the number of queries) is dictated by the target problem at hand, and hence there is no way around this particular barrier. It was pointed out by the author [60] that, at a technical level, another barrier to further progress is perhaps the “many-to-one”-ness of the projection constraints $\pi_e : [m] \mapsto [n]$ in the 2-Prover-1-Round Game instance. As $\delta \rightarrow 0$ in Theorem 2.3, the “many-to-one”-ness $\frac{m}{n} \rightarrow \infty$. In the consistency testing part, one has two hypercubes of widely differing dimensions m and n , and a two query PCP is too weak to enforce consistency between them. This motivated the study of so-called Unique Game problem which is a special case of 2-Prover-1-Round Game problem with $m = n$ and $\pi_e : [m] \mapsto [n]$ are bijections. Such games were considered before in the context of parallel repetition theorem, but the problem of (approximately) computing their value wasn’t considered before. The author conjectured that the problem is computationally hard. We define the Unique Game problem and the Unique Games Conjecture below. Henceforth, \mathcal{U}_{2p1r} will denote a general 2-Prover-1-Round Game, \mathcal{U} (without any sub-script) will denote a Unique Game, and \mathcal{U} along with different sub-scripts will denote different variants of the Unique Game problem .

Definition 2.4 *A Unique Game $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$ is a constraint satisfaction problem defined as follows: $G(V, E)$ is a directed graph whose vertices represent variables and edges represent constraints. The goal is to assign to each vertex a label from the set $[n]$. The constraint on an edge $e = (v, w) \in E$ is described by a bijection $\pi_e : [n] \mapsto [n]$. A labeling $L : V \mapsto [n]$ satisfies the constraint on edge $e = (v, w)$ if and only if $\pi_e(L(v)) = L(w)$. Let $\text{OPT}(\mathcal{U})$ denote the maximum fraction of constraints that can be satisfied by any labeling:*

$$\text{OPT}(\mathcal{U}) := \max_{L:V \mapsto [n]} \frac{1}{|E|} \cdot |\{e \in E \mid L \text{ satisfies } e\}|.$$

The Unique Game problem is to find an (approximately) optimal labeling to a given Unique Game instance. The uniqueness refers to the property that for every edge $e = (v, w)$, fixing a label to v uniquely fixes the label to w and vice versa, since π_e is a bijection. Note that in the definition of 2-Prover-1-Round Game, the constraint graph is bipartite whereas in Definition 2.4, this is not necessarily so. This distinction is minor as can be seen by the following view of the Unique Game problem as a (2-prover and hence bipartite) game: given an instance $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in$

$E\}$ of Unique Game problem, the verifier picks an edge $e = (v, w)$ at random, sends the vertex v to prover P_1 and the vertex w to prover P_2 . Each prover is supposed to answer with a label in $[n]$, and the verifier accepts if and only if $\pi_e(i) = j$ where i and j are answers of the two provers respectively. The strategies of the provers correspond to labelings $L_1, L_2 : V \mapsto [n]$. The *value* of the game is the maximum over all prover strategies, the probability that the verifier accepts and can be easily seen to be between $\text{OPT}(\mathcal{U})$ and $\max\{1, 4 \cdot \text{OPT}(\mathcal{U})\}$. On the other hand, any such game can be thought of as an instance of Unique Game where the underlying graph is bipartite by thinking of questions to the provers as vertices and pairs of questions as edges. This shows that the constraint satisfaction view and the 2-Prover-1-Round view are equivalent.

Let us return to the issue of determining the (approximate) optimum for a Unique Game instance. First note that if $\text{OPT}(\mathcal{U}) = 1$, i.e. when there is a labeling that satisfies every edge, such a labeling can be easily found in polynomial time. This is because, we can assume that we have the correct label for one vertex (by trying out all possible labels to it), and then whenever we know a label for a vertex, it uniquely fixes labels to its neighbors and so on. Thus a labeling can be found for each connected component of the constraint graph.

From the viewpoint of the Unique Games Conjecture, the interesting case is when $\text{OPT}(\mathcal{U}) = 1 - \varepsilon$ for some small positive constant ε . In this case, the above mentioned algorithm to find a good labeling does not seem to work and one may conjecture that finding even a labeling that satisfies a δ fraction of edges is an NP-hard problem.

Conjecture 2.5 Unique Games Conjecture ([60]): *For every $\varepsilon, \delta > 0$, there exists a constant $n = n(\varepsilon, \delta)$, such that given a Unique Game instance $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$, it is NP-hard to distinguish between these two cases:*

- *YES Case:* $\text{OPT}(\mathcal{U}) \geq 1 - \varepsilon$.
- *NO Case:* $\text{OPT}(\mathcal{U}) \leq \delta$.

A few remarks: as we noted, such a conjecture would be false if $\varepsilon = 0$. We note the order of quantifiers: for every ε, δ , there exists a sufficiently large n for which the conjecture holds on Unique Game instances with label set of size n . It is easy to see that a random assignment of labels satisfies $\frac{1}{n}$ fraction of the constraints, and hence for the conjecture to be true, one must necessarily have $n \geq \frac{1}{\delta}$. It is also known that $n \geq 2^{\Omega(1/\varepsilon)}$ for the conjecture to be true (see Section 5). One could consider a formulation with edge-weights, but it is equivalent to the un-weighted case. Also, we can assume by a standard transformation, that the constraint graph is regular, and for the sake of convenience, that the edges are undirected (with implicit understanding of which way the bijection π_e goes).

In terms of the corresponding gap problem, the UGC states that $\text{GapUG}_{1-\varepsilon, \delta}$ is NP-hard for a sufficiently large label set. From the viewpoint of inapproximability, what we really care about is that the gap problem is computationally “hard” (say, it does not have a polynomial time algorithm), and not necessarily that the problem is “NP-hard”. It is possible that the problem is neither in P nor NP-complete, and it would certainly be interesting if there is evidence towards this possibility.

2.5 Boolean Functions, Dictatorships, and Influence of Variables

Let $\mathcal{F} := \{f \mid f : \{-1, 1\}^n \mapsto \{-1, 1\}\}$ be the class of all boolean functions on n -dimensional boolean hypercube. A *balanced* function f takes each of the values $\{-1, 1\}$ equally often, i.e.

$\mathbb{E}[f] = 0$. A $(\frac{1}{3}, \frac{2}{3})$ -balanced function takes each of the values $\{-1, 1\}$ on at most $\frac{2}{3}$ fraction of the inputs.³

A *dictatorship* function $f \in \mathcal{F}$ is defined as a function $f(x) = x_{i_0}$ for some $i_0 \in [n]$. We aim to formalize when a function is considered as *far from being a dictatorship*. Towards this end, let the influence of the i^{th} co-ordinate on a function f be defined as:

$$\text{Inf}_i(f) := \Pr_x [f(x_1, \dots, x_i, \dots, x_n) \neq f(x_1, \dots, -x_i, \dots, x_n)].$$

For a dictatorship function, the relevant co-ordinate has influence 1 and all other influences are zero. Thus we may define a function as far from being a dictatorship if all of its influences are small and this turns out to be a very useful definition. Note that the MAJORITY $:= \text{sign}(\sum_{i=1}^n x_i)$ function has all influences $O(\frac{1}{\sqrt{n}})$ and thus we consider it to be far from a dictatorship.⁴

It is well-known that any function $f : \{-1, 1\}^n \mapsto \mathbb{R}$ has a Fourier expansion:

$$f(x) = \sum_{S \subseteq [n]} \hat{f}(S) \prod_{i \in S} x_i,$$

where the $\hat{f}(S) \in \mathbb{R}$ are the Fourier coefficients. When f is a boolean function, by Parseval's identity, $\sum_S \hat{f}(S)^2 = \mathbb{E}[f^2] = 1$. It is easily proved that:

$$\text{Inf}_i(f) = \sum_{i \in S} \hat{f}(S)^2.$$

We are now ready to informally state a framework in which most of the UGC-based inapproximability results fit. The basic framework was developed by Bellare, Goldreich, and Sudan [18] and then extended in subsequent works. In [18], the Long Code was introduced and it was demonstrated how to combine it with the 2-Prover-1-Round Game problem.⁵ Håstad's [49, 51] work introduced the powerful Fourier methods in analyzing the Long Code based PCPs. The UGC was introduced in [60] and the notion of low influence functions (in PCP context) in [36, 63].

Theorem 2.6 (Informal) *Suppose \mathcal{I} is a maximization problem and $\text{Val} : \mathcal{F} \mapsto \mathbb{R}^+$ is a valuation on boolean functions. Suppose there are constants $0 < s < c$ such that*

1. (Completeness): *If f is a dictatorship, $\text{Val}(f) \geq c$.*
2. (Soundness): *If all influences of f are at most η , then $\text{Val}(f) \leq s$.*

Assume the Unique Games Conjecture (or its variant). Then $\text{Gap}\mathcal{I}_{c,s}$ is NP-hard. In particular, \mathcal{I} is inapproximable within any factor strictly less than c/s .

The theorem is stated in a very informal manner and calls for several comments: (1) The choice of the valuation $\text{Val}(\cdot)$ depends very much on the problem \mathcal{I} and different problems lead to different interesting valuations. Often, $\text{Val}(\cdot)$ is the probability that a function passes a certain dictatorship

³The choice of the balance parameter is rather arbitrary. Any theorem that works for $(\frac{1}{3}, \frac{2}{3})$ -balanced functions, typically also works for $(c, 1 - c)$ -balanced functions for any constant $c > 0$, with possibly a constant factor loss in the parameters.

⁴A random function is intuitively very far from being a dictatorship, but has all influences close to $\frac{1}{2}$ (same holds for the PARITY function whose all influences are 1). This is a somewhat annoying issue and one gets around it by using a more refined notion of “low degree” influence (see [62]). In this article, however, we choose not to present this refined notion for the sake of simplicity.

⁵In PCP literature, this scheme is often referred to as the Outer and Inner PCP Composition.

test. (2) We are interested in the limiting case when $\eta \rightarrow 0$. Often we have $s = s' + \delta$ where s' is a specific constant and $\delta \rightarrow 0$ as $\eta \rightarrow 0$. (3) Sometimes the function needs to be balanced (or roughly balanced) for the soundness condition to make sense. This is often enforced by using a PCP trick known as *folding*, i.e. requiring that functions satisfy $f(-x) = -f(x)$. (4) An analogous theorem holds for minimization problems as well.

The reduction from $\text{GapUG}_{1-\varepsilon, \delta}$ to the gap problem $\text{Gap}\mathcal{I}_{c,s}$ follows the PCP strategy described before. One replaces every vertex of the **Unique Game** instance by the hypercube $\{-1, 1\}^n$, and the PCP proof consists of writing down a boolean function for each of these hypercubes. The valuation $\text{Val}(\cdot)$ dictates the specific form that the codeword testing and consistency testing take. Assume for the sake of illustration that $\text{Val}(\cdot)$ is defined as the probability that a function passes a certain two-query dictatorship test. Then the codeword test consists of running the test on each boolean function separately. The codeword test can be easily extended to a consistency test on two functions $f, g : \{-1, 1\}^n \mapsto \{-1, 1\}$, by reading the first query from the truth-table of f and the second one from that of g . In some sense, the **Unique Games Conjecture** is so powerful that it reduces the analysis of the entire reduction to the analysis of the codeword test (i.e. the gadget) alone.

Clearly, the valuation $\text{Val}(\cdot)$ is supposed to capture a certain property of dictatorship functions. Let us observe a few such properties that are useful:

1. Dictatorships are linear, i.e. $\forall x, y \in \{-1, 1\}^n, f(xy) = f(x)f(y)$, where xy denotes the string that is bitwise product of strings x and y .
2. Dictatorships are stable under noise. Define the ε -noise sensitivity $\text{NS}_\varepsilon(f)$ as $\Pr[f(x) \neq f(y)]$ when $x \in \{-1, 1\}^n$ is chosen uniformly at random and $y \in \{-1, 1\}^n$ is obtained by flipping every bit of x with probability ε . Clearly, if f is a dictatorship, then $\text{NS}_\varepsilon(f) = \varepsilon$. In contrast, **MAJORITY** is less stable and $\text{NS}_\varepsilon(\text{MAJORITY}) = \theta(\sqrt{\varepsilon})$.
3. If $C \subseteq \{-1, 1\}^n$ is a random sub-cube with dimension εn , then with probability $1 - \varepsilon$, a dictatorship function is constant on C . A sub-cube of dimension k is a set of inputs that all agree on a specific setting of input bits outside a set of co-ordinates $T \subseteq [n], |T| = k$.
4. The Fourier mass of a dictatorship function is concentrated at level 1, i.e. on sets $|S| = 1$. In contrast, the Fourier mass of **MAJORITY** at level 1 is very close to $\frac{2}{\pi}$.

2.6 Integrality Gaps

An inapproximability result rules out *every* polynomial time algorithm to compute good approximate solutions. When such results are not available, one could try to rule out specific candidate algorithms. One widely studied class of algorithms is based on linear or semi-definite programming relaxations. Let us consider the **MaxCut** problem as an illustration. Given a graph $G(V = [N], E)$, the goal is to find a partition of V that cuts a maximum fraction of the edges. The SDP relaxation for this problem proposed by Goemans and Williamson [43] appears below:

The **MaxCut** problem can be formulated as the above program with v_i as integer variables restricted to taking values in the set $\{-1, 1\}$. The relaxation allows v_i to be unit vectors in arbitrary dimensional Euclidean space (since there are only N vectors, w.l.o.g. in N -dimensional space). If $\text{OPT}(G)$ denotes the maximum size of a cut in the graph and $\text{SDP}(G)$ denotes the maximum of the relaxation, then clearly $\text{OPT}(G) \leq \text{SDP}(G)$. The “integrality gap” α_{MC} is defined as the least

$\text{Maximize } \frac{1}{ E } \sum_{(i,j) \in E} \frac{1 - v_i \cdot v_j}{2}$ <p style="text-align: left; margin-left: 20px;">Subject to,</p> $\forall i \in [N], \quad \ v_i\ = 1.$ <p style="text-align: center; margin-top: 20px;">Figure 1: SDP Relaxation for MaxCut</p>

upper bound on the ratio $\text{SDP}(G)/\text{OPT}(G)$ over all graphs:

$$\alpha_{\text{MC}} := \sup_G \frac{\text{SDP}(G)}{\text{OPT}(G)}.$$

Clearly, solving the SDP-relaxation gives an approximation to the MaxCut problem within factor α_{MC} . An upper bound of $\alpha_{\text{MC}} \leq 2$ is trivial. Goemans and Williamson [43] proved that:

Theorem 2.7 ([43])

$$\alpha_{\text{MC}} \leq \max_{\theta \in [0, \pi]} \left(\frac{\pi}{2} \frac{1 - \cos \theta}{\theta} \right) = \frac{\pi}{2} \frac{1 - \cos \theta_c}{\theta_c} \approx 1.13,$$

where θ_c is the “critical angle” that achieves the maximum.

Goemans and Williamson [43] give a randomized “rounding algorithm” (by now well-known random hyperplane rounding) that, given any feasible solution to the SDP, finds a cut in the graph whose value is at least $1/\alpha_{\text{MC}}$ times the value of the SDP solution in expected sense. Thus the approximation algorithm not only approximately finds the value of the optimum, but also yields an approximate cut. This feature is shared by almost every algorithm based on LP or SDP relaxation. On the other hand, a lower bound on the integrality gap is taken as evidence that the specific LP/SDP relaxation would not give a better approximation than the lower bound. For the above SDP, Feige and Schechtman [41] show (see also Karloff [57]) a matching lower bound on the integrality gap.

Theorem 2.8 ([41]) *For every $\varepsilon > 0$, there exists a graph $G(V, E)$ such that $\frac{\text{SDP}(G)}{\text{OPT}(G)} \geq \frac{\pi}{2} \frac{1 - \cos \theta_c}{\theta_c} - \varepsilon$. Thus*

$$\alpha_{\text{MC}} = \sup_G \frac{\text{SDP}(G)}{\text{OPT}(G)} \geq \frac{\pi}{2} \frac{1 - \cos \theta_c}{\theta_c},$$

and combined with Theorem 2.7, we get $\alpha_{\text{MC}} = \frac{\pi}{2} \frac{1 - \cos \theta_c}{\theta_c}$.

We note that a lower bound on the integrality gap is a construction of a graph $G(V, E)$ for which the gap is achieved. One needs to prove that $\text{OPT}(G)$ is “low” and $\text{SDP}(G)$ is “high”. The latter is proved by exhibiting a feasible vector solution (not necessarily the one that achieves SDP optimum) that achieves a “high” value. The integrality gap constructions are often explicit, but there are ones using probabilistic arguments. In general, characterizing integrality gaps is a challenging question, and unresolved for many problems, the Grothendieck’s Problem and Sparsest Cut to name a few.

3 Variants of the Unique Games Conjecture

In this Section, we describe several variants of the UGC that are useful to prove inapproximability results. First, we describe a variant that a priori seems much weaker but turns out to be equivalent to the UGC.

Conjecture 3.1 *Weak Unique Games Conjecture:* *There is an increasing unbounded function $\Gamma : \mathbb{R}^+ \mapsto \mathbb{R}^+$ such that the following holds: for every $\varepsilon > 0$, there exists a constant $n = n(\varepsilon)$ such that given a Unique Game instance $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$, it is NP-hard to distinguish between these two cases:*

- *YES Case:* $\text{OPT}(\mathcal{U}) \geq 1 - \varepsilon$.
- *NO Case:* $\text{OPT}(\mathcal{U}) \leq 1 - \sqrt{\varepsilon} \cdot \Gamma(1/\varepsilon)$.

Rao [95] proves a strong parallel repetition theorem which shows that the Conjecture 3.1 implies the Conjecture 2.5 (UGC), see Section 7.

Let $\mathcal{U}_{Lin}(G(V, E), [n], \{\pi_e | e \in E\})$ be a special case of the Unique Game problem where every bijection π_e is of the form $\pi_e(i) = i + c_e \pmod{n} \forall i \in [n]$, for some $c_e \in \mathbb{Z}_n$. One could conjecture that this special case is computationally hard too.

Conjecture 3.2 *For every $\varepsilon, \delta > 0$, there exists a constant $n = n(\varepsilon, \delta)$, such that given a Unique Game instance $\mathcal{U}_{Lin}(G(V, E), [n], \{\pi_e | e \in E\})$, it is NP-hard to distinguish between these two cases:*

- *YES Case:* $\text{OPT}(\mathcal{U}_{Lin}) \geq 1 - \varepsilon$.
- *NO Case:* $\text{OPT}(\mathcal{U}_{Lin}) \leq \delta$.

Khot, Kindler, Mossel, and O’donnell [63] show that Conjecture 3.2 is in fact equivalent to Conjecture 2.5 (UGC). Khot and Regev [68] show that the UGC is equivalent to (an a priori stronger) Conjecture 3.3 below: in the YES case, not only that there is an assignment that satisfies almost all the edges, there is a large subset of vertices such that *all* edges inside it are satisfied.

Conjecture 3.3 *For every $\varepsilon, \delta > 0$, there exists a constant $n = n(\varepsilon, \delta)$, such that given a Unique Game instance $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$, it is NP-hard to distinguish between these two cases:*

- *YES Case:* *There is a subset $S \subseteq V, |S| \geq (1 - \varepsilon)|V|$ such that all edges inside S are satisfied.*
- *NO Case:* $\text{OPT}(\mathcal{U}) \leq \delta$.

As the result of [7] shows, UGC cannot hold when the constraint graph is an expander beyond a certain threshold. One could still hypothesize that the UGC holds when the graph has moderate expansion below this threshold. Interestingly, the threshold depends on the completeness parameter. In the following hypothesis, $t = \frac{1}{2}$ corresponds to this expansion threshold; Arora *et al* [7] show that the hypothesis would be false for $t \leq \frac{1}{2}$.

Hypothesis 3.4 ([7]) *Unique Games Conjecture with Expansion:*⁶ *The following holds for some universal constant $\frac{1}{2} < t < 1$. For every $\varepsilon, \delta > 0$, there exists a constant $n = n(\varepsilon, \delta)$, such that given a Unique Game instance $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$, it is NP-hard to distinguish between these two cases:*

⁶The hypothesis in the conference version [7, Hypothesis 5.1] is stronger and in author’s opinion, probably false. The author tends to believe in the hypothesis stated here; it still implies an inapproximability result for the uniform Sparsest Cut [7, Theorem 5.2]; the proof is to be included in the journal version of [7].

- *YES Case:* $\text{OPT}(\mathcal{U}) \geq 1 - \varepsilon$.
- *NO Case:* $\text{OPT}(\mathcal{U}) \leq \delta$ and moreover, any $(\frac{1}{10}, \frac{9}{10})$ -balanced cut in the graph $G(V, E)$ cuts at least ε^t fraction of its edges.

Bansal and Khot [15] use the following variant of the UGC:

Hypothesis 3.5 ([15]) *Unique Games Conjecture with Vertex Expansion:* For every $\varepsilon, \delta, \gamma > 0$, there exists an integer $n = n(\varepsilon, \delta, \gamma)$ such that for a bipartite Unique Game instance $\mathcal{U}(G(V, W, E), [n], \{\pi_e | e \in E\})$, it is NP-hard to distinguish between:

- *YES Case:* There are sets $V' \subseteq V$, $W' \subseteq W$ such that $|V'| \geq (1 - \varepsilon)|V|$ and $|W'| \geq (1 - \varepsilon)|W|$ and an assignment that satisfies all the edges between V' and W' .
- *NO Case:* $\text{OPT}(\mathcal{U}) \leq \delta$. Moreover, for any two sets $S \subseteq V, T \subseteq W$, $|S| = \gamma|V|, |T| = \gamma|W|$, there is an edge between S and T .

An instance of d -to-1 Game $\mathcal{U}_{d \rightarrow 1}(G(V, W, E), [n], \{\pi_e | e \in E\})$ is a variant where the graph is bipartite with vertex sets V and W , every vertex in V is to receive a label in $[dn]$, every vertex in W is to receive a label in $[n]$, and $\pi_e : [dn] \mapsto [n]$ is a d -to-1 map, i.e. every $j \in [n]$ has exactly d pre-images. For every fixed value of $d \geq 2$, one could conjecture that finding good labelings to d -to-1 games is hard.

Conjecture 3.6 ([60]) *$d \mapsto 1$ Conjecture:* For every $\delta > 0$, there exists a constant $n = n(\delta)$, such that given an instance $\mathcal{U}_{d \rightarrow 1}(G(V, W, E), [dn], [n], \{\pi_e | e \in E\})$, it is NP-hard to distinguish between these two cases:

- *YES Case:* $\text{OPT}(\mathcal{U}_{d \rightarrow 1}) = 1$.
- *NO Case:* $\text{OPT}(\mathcal{U}_{d \rightarrow 1}) \leq \delta$.

Note that in the d -to-1 Conjecture, one may hope that the instance is satisfiable in the YES Case. An instance of \times -Game $\mathcal{U}_{\times}(G(V, E), [n], \{(\sigma_e, \rho_e) | e \in E\})$ is a variant where n is even and the constraint on edge $e = (v, w)$ consists of two permutations $\sigma_e, \rho_e : [n] \mapsto [n]$, and a labeling $(L(v), L(w))$ satisfies the edge if and only if

$$(L(v), L(w)) \in \left\{ (\sigma_e(2i-1), \rho_e(2i-1)), (\sigma_e(2i-1), \rho_e(2i)), (\sigma_e(2i), \rho_e(2i-1)) \right\} \text{ for some } 1 \leq i \leq \frac{n}{2}.$$

Dinur, Mossel, and Regev [35] make the following conjecture:⁷

Conjecture 3.7 ([35]) *\times -Conjecture:* For every $\delta > 0$, there exists a constant $n = n(\delta)$, such that given an instance $\mathcal{U}_{\times}(G(V, E), [n], \{(\sigma_e, \rho_e) | e \in E\})$, it is NP-hard to distinguish between these two cases:

- *YES Case:* $\text{OPT}(\mathcal{U}_{\times}) = 1$.
- *NO Case:* $\text{OPT}(\mathcal{U}_{\times}) \leq \delta$.

⁷The conjecture in [35] is somewhat weaker. We take the liberty to state a revised conjecture here.

4 Inapproximability Results

The main motivation behind formulating the UGC is to prove inapproximability results for NP-complete problems for which “satisfactory” results are not available otherwise. Figure 2 shows a table of such results. For each problem, we mention the best approximation factor known, the inapproximability factor under the UGC (or its variant), and the best inapproximability result known without relying on the UGC. The list is not necessarily exhaustive, but covers most of the results the author is aware of. The problems can be roughly classified into two categories: problems for which a constant approximation is known and the goal is to characterize the exact approximation threshold (e.g. MaxCut), and those for which the known approximation is super-constant, and the goal is to prove a super-constant inapproximability result (e.g. Min-2SAT-Deletion). Majority of these inapproximability results can be proved in the framework described in Theorem 2.6. We present a complete reduction from GapUG to GapMin-2SAT-Deletion as an illustration. We refer to the author’s recent article [62] for a more detailed survey of such reductions.

4.1 Inapproximability of the Min-2SAT-Deletion Problem

In this section, we give a sketch of a reduction from the Unique Game problem to the Min-2SAT-Deletion problem [60] as one illustration of the UGC-based inapproximability results. An instance of the Min-2SAT-Deletion problem is a 2SAT formula ϕ and $\text{OPT}(\phi)$ is defined as the minimum fraction of clauses that need to be deleted from ϕ to make it satisfiable.

Theorem 4.1 *Let $\frac{1}{2} < c < 1$ be a fixed constant. Then, for every sufficiently small constant $\varepsilon > 0$, there exists a constant $\delta > 0$ and a polynomial time reduction that maps an instance $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$ of the Unique Game problem to a 2SAT formula ϕ such that:*

- (YES Case:) If $\text{OPT}(\mathcal{U}) \geq 1 - \varepsilon$, then $\text{OPT}(\phi) \leq O(\varepsilon)$.
- (NO Case:) If $\text{OPT}(\mathcal{U}) \leq \delta$, then $\text{OPT}(\phi) \geq \Omega(\varepsilon^c)$.

We actually give a reduction from the Unique Game problem to a boolean 2-CSP where the constraints are of the form $x = y$ for boolean literals x and y . The constraints can then be converted to 2SAT clauses since $x = y$ is equivalent to two clauses $x \vee \bar{y}$ and $\bar{x} \vee y$. It is easier to view the reduction as a PCP where the verifier reads two bits from a proof and accepts if and only if the two bits are equal (or unequal, as is pre-determined). The 2-CSP can then be written down by letting the bits in the proof as boolean variables and the tests of the verifier as the constraints of the 2-CSP. In the YES Case, there is a proof that the verifier accepts with probability $1 - O(\varepsilon)$ whereas in the NO Case, every proof is accepted with probability at most $1 - \Omega(\varepsilon^c)$. In terms of the 2-CSP, this means that in the YES Case, it suffices to delete $O(\varepsilon)$ fraction of the constraints to make the 2-CSP satisfiable, and in the NO Case, one needs to delete at least $\Omega(\varepsilon^c)$ fraction of the constraints. An important ingredient in the PCP analysis is the Bourgain’s Noise-Sensitivity Theorem that is stated in Section 6.2 as Theorem 6.2.

The PCP construction is rather straightforward. The verifier is given a Unique Game instance $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$. She expects the proof to contain, for every $v \in V$, the long code of the label of v in a supposed labeling $L : V \mapsto [n]$. In words, for every $v \in V$, if $L(v) = i_0$, then the proof is supposed to contain the truth-table of the dictatorship function of co-ordinate i_0 . Each of these functions is assumed to satisfy $f(-x) = -f(x)$ and hence balanced. This is enforced by the standard “folding” trick: we force the proof to contain only one of the bits $f(x)$ or $f(-x)$ and let the other bit to be its negation.

Problem	Best Approx. Known	Inapprox. Known Under UGC	Best Inapprox. Known	Ref.
Vertex Cover (VC)	2	$2 - \varepsilon$	1.36	[68, 15, 36]
VC on k -uniform Hypergraphs, $k \geq 3$	k	$k - \varepsilon$	$k - 1 - \varepsilon$	[68, 16, 34]
MaxCut	α_{MC}	$\alpha_{MC} - \varepsilon$	$\frac{17}{16} - \varepsilon$	[43, 63, 51] [87, 66]
Max-2SAT*	α_{LLZ}	$\alpha_{LLZ} - \varepsilon$	APX-hard*	[78, 12]
Any CSP \mathcal{C} with integrality gap $\alpha_{\mathcal{C}}$	$\alpha_{\mathcal{C}}$	$\alpha_{\mathcal{C}} - \varepsilon$		[90]
Max- k CSP	$O(2^k/k)$	$\Omega(2^k/k)$	$2^{k-O(\sqrt{k})}$	[23, 101, 13, 100]
Max-3CSP on satisfiable instances	$\frac{8}{5}$	$\frac{8}{5} - \varepsilon$, under Conj. 3.6	$\frac{27}{20} - \varepsilon$	[105, 88, 69]
Max Acyclic Subgraph	2	$2 - \varepsilon$	$\frac{66}{65} - \varepsilon$	[48, 86]
Feedback Arc Set	$\tilde{O}(\log N)$	$\omega(1)$	APX-hard	[48, 102]
Non-uni. Sparsest Cut	$\tilde{O}(\sqrt{\log N})$	$\omega(1)$	APX-hard	[8, 25, 71, 31]
Uniform Sparsest Cut,	$O(\sqrt{\log N})$	$\omega(1)$, under Hypo. 3.4	No PTAS	[10, 7, 3]
Min-2SAT-Deletion, Min-Uncut	$O(\sqrt{\log N})$	$\omega(1)$	APX-hard	[1, 60]
Coloring 3-colorable Graphs	$N^{.2111}$	$\omega(1)$, under Conj. 3.7	5	[6, 35, 58]
Coloring $2d$ -colorable Graphs, $d \geq 2$	$N^{1-\frac{3}{2d+1}}$	$\omega(1)$, under Conj. 3.6	$2d + 2 \lfloor \frac{2d}{3} \rfloor - 1$, $d^{\Omega(\log d)}$	[55, 35] [58, 59]
Scheduling with Prec. Constraints*	2	$2 - \varepsilon$, under Hypo. 3.5		[15]
Kernel Clustering kernel matrix B	$K(B)$	$K(B) - \varepsilon$	APX-hard	[64, 65]
L_p Grothendieck Prob.* $p > 2$	γ_p^2	$\gamma_p^2 - \varepsilon$		[72] + Follow-up.
Multiway Cut integrality gap $\alpha \leq 1.344$	α	$\alpha - \varepsilon$	APX-hard	[21, 56, 81]

Figure 2: (In)approximability Results for Some Optimization Problems

* α_{LLZ} is the believed approximation ratio of Lewin *et al*'s SDP algorithm [78].

* APX-hard means NP-hard to approximate within some factor strictly larger than 1.

* Minimizing weighted completion time with precedence constraints.

* $K(B)$ is a geometric parameter associated with matrix B .

* γ_p is the p^{th} norm of a standard Gaussian variable.

The PCP Test

- Pick a random edge $(v, w) \in E$ along with the bijection $\pi = \pi_e : [n] \mapsto [n]$ associated with it. Let f_v and f_w be the supposed long codes corresponding to v and w respectively.
- Let $x \in \{-1, 1\}^n$ be a random input and $y \in \{-1, 1\}^n$ be obtained by flipping every bit of x with probability $\varepsilon > 0$ independently.
- Do each of these two tests with probability $\frac{1}{2}$ each: accept if $f_v(x) = f_v(y)$ or $f_v(x) = g_w(\pi(x))$, where $\pi(x) := (x_{\pi(1)}, \dots, x_{\pi(n)})$.

YES Case: Assume the Unique Game instance has a labeling that satisfies $1 - \varepsilon$ fraction of the edges. Take this labeling and encode each label with the corresponding dictatorship function. Since the ε -noise sensitivity of a dictatorship function is exactly ε , the test $f_v(x) = f_v(y)$ accepts with probability $1 - \varepsilon$. For $1 - \varepsilon$ fraction of the edges, the labeling satisfies the edge, i.e. f_v and f_w are dictatorships of co-ordinates i and j respectively with $\pi(i) = j$. Hence $f_v(x) = x_i$ and $g_w(\pi(x)) = (\pi(x))_j = x_{\pi(j)} = x_i$, and hence the test $f_v(x) = g_w(\pi(x))$ accepts with probability 1. Overall, the verifier accepts with probability $1 - O(\varepsilon)$.

NO Case: Assume for the sake of contradiction that the verifier accepts with probability at least $1 - \frac{\varepsilon^c}{200}$. We will construct a labeling to the Unique Game instance that satisfies more than a δ fraction of its edges, giving a contradiction. Towards this end, first observe that by an averaging argument, it must be the case that for 99% of the vertices $v \in V$, the test $f_v(x) = f_v(y)$ accepts with probability at least $1 - \varepsilon^c$. Since f_v is a balanced function, applying Bourgain's Theorem 6.2, we conclude that f_v is *close* to a function that depends only on $K := 2^{O(1/\varepsilon^2)}$ co-ordinates. Let this set of co-ordinates be $S_v \subseteq [n], |S_v| \leq K$. For the ease of argument, assume that this holds for every $v \in V$ and moreover that f_v depends precisely on the co-ordinates in S_v . Now consider an edge $(v, w) \in E$. We claim that for 99% of edges, the sets $\pi(S_v)$ and S_w intersect. This is because, otherwise, the functions f_v and f_w depend on disjoint sets of co-ordinates (after the bijection π is applied) and therefore the test $f_v(x) = f_w(\pi(x))$ fails with constant probability. Finally, consider a labeling that assigns to each $v \in V$, a random label from the set S_v . It follows that for 99% of the edges $(v, w) \in E$, the sets $\pi(S_v)$ and S_w intersect and hence a matching pair of labels is assigned to v and w with probability at least $\frac{1}{K}$. We choose $\delta \ll \frac{1}{K}$ beforehand to reach a contradiction.

4.2 Raghavendra's Result

A beautiful result of Raghavendra [90] shows that for every CSP \mathcal{C} , the integrality gap $\alpha_{\mathcal{C}}$ for a natural SDP relaxation is same as the inapproximability threshold for the CSP, modulo the Unique Games Conjecture. This is a duality type result showing that the SDP relaxation, while giving an algorithm with approximation ratio $\alpha_{\mathcal{C}}$, also leads to an inherent limitation: it prevents *every* efficient algorithm from achieving a strictly better approximation ratio.

A k -ary CSP \mathcal{C} is defined by a collection of finitely many predicates $P : \{0, 1\}^k \mapsto \{0, 1\}$. By abuse of notation, we let \mathcal{C} denote this collection of predicates as well. An instance of the CSP is specified as $I(V, E, \{P_e | e \in E\})$: $V = \{x_1, \dots, x_N\}$ is a set of boolean variables and E is a collection of constraints, each being a size k subset of V . A constraint $e \in E$ is denoted as $e = (x_{e_1}, \dots, x_{e_k}), e_i \in [N]$, with a specific order on the variables, and has an associated predicate $P_e \in \mathcal{C}$. An assignment is a map $\rho : V \mapsto \{0, 1\}$. An assignment ρ satisfies a constraint e if

$$P_e(\rho(x_{e_1}), \dots, \rho(x_{e_k})) = 1.$$

Let $\text{OPT}(I)$ denote the maximum fraction of constraints satisfied by any assignment. Raghavendra studies a natural SDP relaxation for the CSP that we do not describe here. Let $\text{SDP}(I)$ denote the optimum of the SDP relaxation, and define the integrality gap α_C as:

$$\alpha_C := \sup_I \frac{\text{SDP}(I)}{\text{OPT}(I)},$$

where the supremum is taken over all instances. Clearly, $\text{OPT}(I) \leq \text{SDP}(I) \leq \alpha_C \cdot \text{OPT}(I)$, and hence the SDP gives an α_C approximation to the CSP.⁸ The formal statement of Raghavendra’s result is:

Theorem 4.2 ([90]) *Suppose there is an instance I^* of the CSP \mathcal{C} such that $\text{SDP}(I^*) \geq c$ and $\text{OPT}(I^*) \leq s$. Then for every $\gamma > 0$, there exist $\varepsilon, \delta > 0$, and a polynomial time reduction from an Unique Game instance to an instance I of the CSP such that:*

- (YES Case:) *If $\text{OPT}(U) \geq 1 - \varepsilon$, then $\text{OPT}(I) \geq c - \gamma$.*
- (NO Case:) *If $\text{OPT}(U) \leq \delta$, then $\text{OPT}(I) \leq s + \gamma$.*

In particular, assuming the UGC, it is NP-hard to approximate the CSP within any factor strictly less than α_C .

Note that the theorem turns an integrality gap to an inapproximability gap. Roughly speaking, the idea is to use the SDP gap instance to construct a dictatorship test and then combine the test (as usual) with the Unique Game instance. The analysis of the dictatorship test relies crucially on a powerful version of the *invariance principle* by Mossel [82] (see Section 6.4 for a more detailed description of the invariance principle).

The result of Khot *et al* [63] shows that assuming the UGC, it is NP-hard to approximate MaxCut within any factor strictly less than α_{MC} , the integrality gap of the Goemans-Williamson’s SDP. The result, in hindsight, can be interpreted as turning the Feige-Schechtman integrality gap instance into an inapproximability result. However, the result relied on explicitly “knowing” the integrality gap instance. Raghavendra’s result is a generalization of [63] to every CSP. For a general CSP, the integrality gap is not explicitly known, but Raghavendra circumvents the need to know it explicitly.

Raghavendra’s result is more general: it applies to CSPs with non-boolean alphabet, with non-negative real valued “pay-off functions” rather than predicates, and in some restricted cases, even with possibly negative pay-offs. As a concrete example, Raghavendra and Steurer [93] show that for the Grothendieck’s Problem, even though the integrality gap K_G (= the famous Grothendieck’s constant) is not known, we do know, modulo UGC, that it is NP-hard to approximate Grothendieck’s Problem within any factor strictly less than K_G . Manokaran *et al* [81] show how to translate linear programming integrality gap to an inapproximability gap, specifically for the Multiway Cut and related problems.

5 Algorithms

The UGC states that given an instance of the Unique Game problem that is $1 - \varepsilon$ satisfiable, it is NP-hard to find an assignment that satisfies δ fraction of the constraints when the number of

⁸The SDP gives the approximate *value* of the solution. Raghavendra and Steurer [91] also show how to round the solution and find an actual assignment with approximation ratio arbitrarily close to α_C .

Algorithm	Value of Solution Found on $1 - \varepsilon$ satisfiable instance with n labels and N vertices
Khot [60]	$1 - O(n^2 \varepsilon^{1/5} \sqrt{\log(1/\varepsilon)})$
Trevisan [103]	$1 - O(\sqrt[3]{\varepsilon \log N})$
Gupta, Talwar [46]	$1 - O(\varepsilon \log N)$
Charikar, Makarychev, Makarychev [22]	$1 - O(\sqrt{\varepsilon \log n})$ $n^{-\varepsilon/2}$
Chlamtac, Makarychev, Makarychev [30]	$1 - O(\varepsilon \sqrt{\log N \log n})$
Arora <i>et al</i> [7]	$1 - O(\varepsilon \cdot \frac{1}{\lambda} \log(\frac{\lambda}{\varepsilon}))$ on graphs with eigenvalue gap λ .
Arora, Barak, Steurer [5]	$1 - \varepsilon^\alpha$ for some $0 < \alpha < 1$ in time $\exp(N^{\varepsilon^\alpha})$.

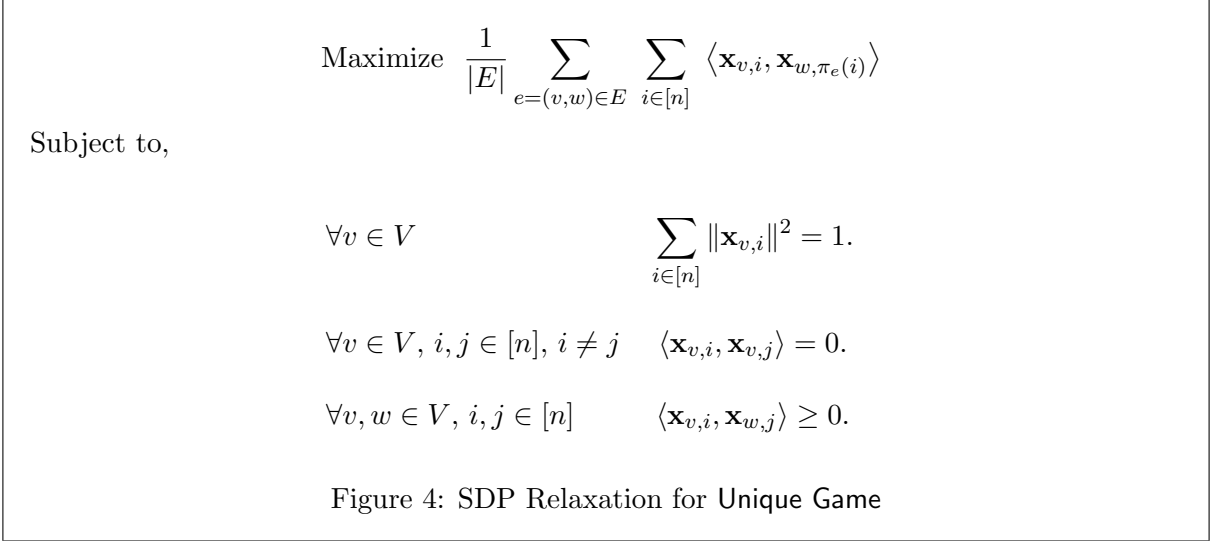
Figure 3: Algorithms for the Unique Game Problem

labels $n = n(\varepsilon, \delta)$ is a sufficiently large constant. Therefore a *good enough* algorithm for the Unique Game problem would disprove the conjecture. Figure 3 summarizes the algorithmic results so far: given an $1 - \varepsilon$ satisfiable instance with N vertices and n labels, the second column specifies the fraction of the constraints satisfied by an assignment found by the algorithm.

We note that the algorithmic results fall short of disproving the UGC: In [60, 22], for constant $\varepsilon > 0$, the algorithm does not work when n is sufficiently large but still constant (which the conjecture allows). In [103, 46, 30], the algorithm works only when ε is a sub-constant function of the instance size N whereas the conjecture is about constant $\varepsilon > 0$. In [7], the algorithm works only when the graph is a (mild) expander. In [5], the algorithm runs in sub-exponential time. On the other hand, these algorithms can be interpreted as meaningful trade-offs between quantitative parameters if the UGC were true: one must have $n \gg 2^{1/\varepsilon}$ [22], the graph cannot even be a mild expander and in particular its eigenvalue gap must satisfy $\lambda \ll \varepsilon$ [7], and the presumed reduction from 3SAT to the GapUG problem must blow up the instance size by a polynomial factor where the exponent of the polynomial grows as $\varepsilon \rightarrow 0$ [5]. Moreover, the conjecture cannot hold for $\varepsilon \ll \frac{1}{\sqrt{\log N}}$ [103, 46, 30]; this places a limitation on the inapproximability factor that can be proved for problems such as Sparsest Cut, where one is interested in the inapproximability factor as a function of the input size.

We present a very brief overview of some of these algorithms. With the exception of [46, 5], all algorithms depend on a natural SDP relaxation for the Unique Game problem shown in Figure 4. We skip the minor differences in the SDP formulations used in different papers. The SDP relaxation was introduced and studied by Feige and Lovasz [39] in the context of parallel repetition (see Section 7).

For a Unique Game instance $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$, one can write a quadratic integer program in the following manner: for every $v \in V$ and $i \in [n]$, let $x_{v,i}$ be a $\{0, 1\}$ -valued integer variable such that $x_{v,i} = 1$ if and only if $L(v) = i$ for some supposed labeling $L : V \mapsto [n]$. Clearly one can write down the constraints that $\sum_{i=1}^n x_{v,i}^2 = 1$, $x_{v,i}x_{v,j} = 0$ for $i \neq j$, and that $x_{v,i}x_{w,j} \geq 0$.



The fraction of constraints satisfied is:

$$\frac{1}{|E|} \sum_{e=(v,w) \in E} \sum_{i \in [n]} x_{v,i} \cdot x_{w,\pi_e(i)}.$$

The SDP relaxation in Figure 4 is obtained by relaxing integer variables $x_{v,i}$ to vectors $\mathbf{x}_{v,i}$ and replacing integer products by vector inner products.

For simplicity of exposition, let us restrict ourselves to the case of linear unique games, as in the definition preceding Conjecture 3.2. Such games are symmetric in the following sense: if $L : V \mapsto [n]$ is a labeling, then for any constant $c \in \mathbb{Z}_n$, the labeling $L'(v) = L(v) + c \pmod n$ is another labeling that satisfies exactly the same set of edges. This symmetry allows us to add more constraints to the SDP and ensure that in the SDP solution, $\forall v \in V, \forall i \in [n], \|\mathbf{x}_{v,i}\| = \frac{1}{\sqrt{n}}$. Thus, for every $v \in V$, we have an orthonormal tuple of vectors $\{\mathbf{x}_{v,i}\}_{i=1}^n$ up to a normalizing factor of $\frac{1}{\sqrt{n}}$. Assume that we have a Unique Game instance that is $1 - \varepsilon$ satisfiable. Thus, the SDP optimum is also at least $1 - \varepsilon$, and it is an average over the contribution over all edges. For simplicity of exposition, let us also assume that the contribution over every edge $e = (v, w)$ is at least $1 - \varepsilon$. For linear games, by symmetry, this ensures that $\forall i \in [n]$,

$$\langle \mathbf{x}_{v,i}, \mathbf{x}_{w,\pi_e(i)} \rangle \geq (1 - \varepsilon) \cdot \frac{1}{n}.$$

5.1 Khot's [60] and Charikar, Makarychev, Makarychev's [22] Algorithms

Assume that the vectors are in \mathbb{R}^d and consider the following natural rounding procedure:

SDP Rounding:

- Pick a random unit vector \mathbf{r} in \mathbb{R}^d .
- For each $v \in V$, define its label to be $L(v) = i_0$ where

$$i_0 = \arg \max_{i \in [n]} |\langle \mathbf{r}, \mathbf{x}_{v,i} \rangle|.$$

The analysis works on edge by edge basis. Fix any edge $e = (v, w)$. We have two tuples $\{\mathbf{x}_{v,i}\}_{i=1}^n$ and $\{\mathbf{x}_{w,\pi_e(i)}\}_{i=1}^n$ for vertices v and w respectively. As we said before, we will make simplifying assumptions that $\forall i, j, \|\mathbf{x}_{v,i}\| = \|\mathbf{x}_{w,j}\| = \frac{1}{\sqrt{n}}$ and that $\langle \mathbf{x}_{v,i}, \mathbf{x}_{w,\pi_e(i)} \rangle \geq (1 - \varepsilon) \cdot \frac{1}{n}$. Thus the two tuples are very “close” and it is intuitively clear that for a random vector \mathbf{r} , the vectors in the two tuples that maximize the inner product with \mathbf{r} are likely to be a “matching” pair of vectors. Formally, we need to analyze the probability that the rounding algorithm “succeeds” on the edge, i.e. the probability that it assigns a label i_0 to v and the matching label $\pi_e(i_0)$ to w . In [60], a rather crude estimate is given, whereas Charikar, Makarychev, and Makarychev [22] give the optimal $1 - O(\sqrt{\varepsilon \log n})$ bound on this probability. For a general unique game (i.e. not necessarily linear), in an SDP solution, the vectors $\{\mathbf{x}_{v,i}\}_{i=1}^n$ need not have equal norms. In this case, the rounding procedure in [22] is considerably more sophisticated and the analysis technically challenging. A rough intuition is that $\{\|\mathbf{x}_{v,i}\|^2\}_{i=1}^n$ can be thought of as a probability distribution on the labels for a vertex v and the rounding procedure should pick a label i with probability $\|\mathbf{x}_{v,i}\|^2$, simultaneously for all vertices in a coordinated manner. The result is applicable when $\varepsilon \ll \frac{1}{\log n}$. In the regime when $\varepsilon \gg \frac{1}{\log n}$, the authors of [22] give an alternate rounding procedure that satisfies $n^{-\varepsilon/2}$ fraction of the constraints.

5.2 Trevisan’s Algorithm [103]

Trevisan [103] presents a simple and elegant algorithm that first partitions the Unique Game graph into components with low diameter and then on each component, finds a good labeling. The SDP relaxation is augmented with the so-called triangle inequality constraints that we omit from this discussion. Assume, again for simplicity, that the Unique Game instance is $1 - \varepsilon$ satisfiable and every edge contributes at least $1 - \varepsilon$ towards the SDP objective. We first apply a graph decomposition theorem of Leighton and Rao [77] that removes ε fraction of the edges so that each connected component has diameter at most $O(\log N/\varepsilon)$. We “give up” on the edges that are removed and find a good labeling for each of the components separately. On each component, fix one vertex v and consider the vectors $\{\mathbf{x}_{v,i}\}_{i=1}^n$. Choose the label of v to be $i_0 \in [n]$ with probability $\|\mathbf{x}_{v,i_0}\|^2$. Once the label i_0 for v is fixed, for every other vertex w in the component, define its label to be $j \in [n]$ that minimizes $\|\mathbf{x}_{w,j} - \mathbf{x}_{v,i_0}\|^2$. Intuitively, since the component has bounded diameter, the label at one vertex more or less determines labels of all vertices in the component (e.g. consider the extreme case when a component is a clique. In this case, the vector tuples for all the vertices in the component are pairwise close to each other). The actual analysis is fairly simple using the triangle inequality constraints.

5.3 Arora *et al*’s Algorithm [7]

This algorithm is again based on the SDP relaxation and works as long as the graph is a mild expander. Let us again make the simplifying assumptions that $\forall v \in V, \forall i, \|\mathbf{x}_{v,i}\| = \frac{1}{\sqrt{n}}$ and that for every edge $e = (v, w) \in E$, $\langle \mathbf{x}_{v,i}, \mathbf{x}_{w,\pi_e(i)} \rangle \geq (1 - \varepsilon) \cdot \frac{1}{n}$. A useful observation is that one can define a single unit vector $\mathbf{x}_v := n^{3/2} \cdot \sum_{i=1}^n \mathbf{x}_{v,i}^{\otimes 4}$ that captures the behavior of the entire tuple in the following sense: for every edge (v, w) , we know that the tuples $\{\mathbf{x}_{v,i}\}_{i=1}^n$ and $\{\mathbf{x}_{w,\pi_e(i)}\}_{i=1}^n$ are close and therefore it holds that $\langle \mathbf{x}_v, \mathbf{x}_w \rangle \geq 1 - O(\varepsilon)$. Now consider the random hyperplane cut on the set of vectors $S := \{\mathbf{x}_v \mid v \in V\}$. Each edge (v, w) is cut with probability $O(\sqrt{\varepsilon})$ since $\langle \mathbf{x}_v, \mathbf{x}_w \rangle \geq 1 - O(\varepsilon)$. If the set S were “well-spread” on the unit sphere, then the cut will be roughly balanced. This would imply that the Unique Game graph has a roughly balanced cut that cuts $O(\sqrt{\varepsilon})$ fraction of the constraints, contradicting the assumption that the graph is a mild expander. Therefore, the set S cannot be well-spread, i.e. most vectors in S are confined to a small

neighborhood on the unit sphere. This implies, in turn, that the tuples $\{\mathbf{x}_{v,i}\}_{i=1}^n$ for almost all $v \in V$ are pairwise close to each other (we already know that the tuples are close whenever there is an edge/constraint between them, but now this holds irrespective of it) and this easily yields a good global labeling. The expansion requirement in [7] is stated in terms of eigenvalue gap, but as we indicated, it also works w.r.t. edge expansion. This was also observed and the result somewhat improved in [80].

5.4 Arora, Barak, Steurer’s Algorithm [5]

For some $0 < \alpha < 1$, Arora, Barak, and Steurer [5] give an algorithm that runs in time $\exp(N^{\varepsilon^\alpha})$ and finds a labeling that satisfies $1 - \varepsilon^\alpha$ fraction of the edges. The algorithm is motivated by the connection between **Unique Game** problem and the **Small Set Expansion** problem mentioned below. The algorithm is based on the approach of Naor [84] and given a graph $G(V, E)$ whose stochastic matrix has at least M eigenvalues larger than $1 - \varepsilon$, the algorithm finds a set $S \subseteq V$ of size $N^{\text{poly}(\varepsilon)}/M$ with expansion at most $\text{poly}(\varepsilon)$. Ideas of Kolla [74] are also used who gives a method for solving unique games on graphs with few large eigenvalues. The author just received the manuscript of [5] at the time of writing this article, so prefers to include only this brief description.

5.5 The Small Set Expansion Problem:

Raghavendra and Steurer [94] give a reduction from the **Small Set Expansion** problem to the **Unique Game** problem. Whereas we have several reductions from the **Unique Game** problem to other optimization problems, this is the first reduction in the other direction. It would indeed be interesting if some natural problem is equivalent to the **Unique Game** problem, and hence “UG-complete”. The **Small Set Expansion** problem is certainly a candidate and so is the $\text{GapMaxCut}_{1-\varepsilon, 1-b\sqrt{\varepsilon}}$ problem as discussed in Section 7.

For a d -regular graph $G(V, E)$ and a set $S \subseteq V$, define the expansion of the set S as $\phi(S) := \frac{|E(S, V \setminus S)|}{d \cdot |S|}$, i.e. the fraction of edges incident on S that leave S . One could hypothesize that:

Hypothesis 5.1 *For every $\varepsilon > 0$, there exists $\gamma > 0$ such that given a regular graph $G(V, E)$, it is NP-hard to distinguish between:*

- (YES Case:) *There is a set $S \subseteq V, |S| = \gamma|V|$ such that $\phi(S) \leq \varepsilon$.*
- (NO Case:) *For every set $S \subseteq V, |S| = \gamma|V|$, $\phi(S) \geq 1 - \varepsilon$.*

In words, it is hard to find a small set in a graph that is somewhat non-expanding even when we are guaranteed that there exists a small set that is almost non-expanding. Raghavendra and Steurer [94] show that this conjecture implies the **Unique Games Conjecture**.

Theorem 5.2 ([94]) *Hypothesis 5.1 implies Conjecture 2.5 (UGC).*

We give some intuition here why the UGC is related to the small set expansion problem. Suppose $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$ is an instance of **Unique Game**. Consider the following graph $G'(V', E')$:

$$V' := V \times [n] \quad \text{and} \quad ((v, i), (w, j)) \in E' \text{ if and only if } (v, w) = e \in E \text{ and } \pi_e(i) = j.$$

In words, G' is obtained by replacing every vertex of G by a block of n vertices and every bijective constraint of the **Unique Game** instance is replaced by a matching between the two blocks. It is easy to see that if the **Unique Game** instance has a labeling $L : V \mapsto [n]$ that satisfies $1 - \varepsilon$ fraction of the constraints, then the set $S := \{(v, L(v)) | v \in V\} \subseteq V'$ is a subset with relative size $\frac{1}{n}$ and expansion at most ε . Thus a good labeling to the **Unique Game** instance corresponds to a small non-expanding set in G' .

6 Discrete Fourier Analysis

In this Section, we give a brief overview of Fourier analytic theorems that have found applications to inapproximability. We refer to [62] for a more detailed exposition.

6.1 Friedgut’s Theorem

Theorem 6.1 ([42]) *Let $k, \gamma > 0$ and let $f : \{-1, 1\}^n \mapsto \{-1, 1\}$ be a boolean function with total influence (a.k.a. average sensitivity) at most k . Then there exists a boolean function g that agrees with f on $1 - \gamma$ fraction of inputs and depends on at most $2^{O(k/\gamma)}$ co-ordinates.*

Friedgut’s Theorem was used by Dinur and Safra [36] to prove 1.36 inapproximability result for the Vertex Cover problem. Their paper, and to a lesser extent the application of Bourgain’s Theorem in [60], were the early examples where a “special purpose” powerful Fourier analytic theorem was applied in the context of inapproximability. Friedgut’s Theorem was subsequently used to prove UGC-based $2 - \varepsilon$ inapproximability result for Vertex Cover [68]. A theorem of Russo [99] is also used in the Vertex Cover applications: for a *monotone* boolean function f , the probability $\Pr[f = 1]$ under a p -biased distribution is an increasing, differentiable function of p , and its derivative equals the average sensitivity.⁹

6.2 Bourgain’s Noise-Sensitivity Theorem

Recall that for $0 < \varepsilon < \frac{1}{2}$, the ε -noise sensitivity $\text{NS}_\varepsilon(f)$ is defined as the probability $\Pr[f(x) \neq f(y)]$ where x is a uniformly random input and the input y is obtained by flipping every bit of x independently with probability ε .

Theorem 6.2 ([20]) *Let $\frac{1}{2} < c < 1$ be any fixed constant. For all sufficiently small $\varepsilon > 0$, if $f : \{-1, 1\}^n \mapsto \{-1, 1\}$ is a function with ε -noise sensitivity at most ε^c , then there is a boolean function g that agrees with f on 99% of the inputs and g depends only on $2^{O(1/\varepsilon^2)}$ co-ordinates.*

Bourgain’s Theorem is tight: for $c = \frac{1}{2}$, the MAJORITY function has ε -noise sensitivity $\theta(\varepsilon^c)$, but cannot be approximated by a function that depends on a bounded number of co-ordinates. The theorem was motivated by Håstad’s [50] observation that such a theorem could lead to a super-constant inapproximability result for the Min-2SAT-Deletion problem. It was observed in [60] that the desired inapproximability result can indeed be proved, albeit using the Unique Games Conjecture. The theorem was subsequently used to prove a super-constant inapproximability result as well as $(\log \log N)^{1/7}$ integrality gap result for the non-uniform Sparsest Cut problem [71]. The gap holds for a natural SDP with so-called *triangle inequality constraints* and can be interpreted as a construction of an N -point *negative type* metric that incurs distortion $(\log \log N)^{1/7}$ to embed into ℓ_1 . It refuted a conjecture of Goemans and Linial [44, 79] that every negative type metric embeds into ℓ_1 with constant distortion (which would have implied a constant factor approximation to non-uniform Sparsest Cut). These connections are further explained in Section 8.2.

6.3 Kahn, Kalai, Linal’s Theorem

Theorem 6.3 ([54]) *Every $(\frac{1}{3}, \frac{2}{3})$ -balanced boolean function $f : \{-1, 1\}^n \mapsto \{-1, 1\}$ has a variable whose influence is $\Omega\left(\frac{\log n}{n}\right)$.*

⁹ f is monotone if flipping any bit of an input from -1 to 1 can only change the value of the function from -1 to 1 . A p -biased distribution is a product distribution such that on each co-ordinate, $\Pr[x_i = 1] = p$.

The KKL Theorem has been used to prove a super-constant inapproximability result for the non-uniform Sparsest Cut problem by Chawla *et al* [25] and to construct $\Omega(\log \log N)$ integrality gaps for a natural SDP relaxation for the non-uniform as well as the uniform Sparsest Cut problem [75, 33], improving on the integrality gap obtained earlier in [71].

6.4 Majority Is Stablest and Borell’s Theorem

Bourgain’s Theorem gives a lower bound of $\Omega(\varepsilon^c)$ on the noise sensitivity of a balanced function whose all influences are small and $c > \frac{1}{2}$. The Majority Is Stablest Theorem [83] gives an exact lower bound, namely $\frac{1}{\pi} \arccos(1 - 2\varepsilon)$, which coincides with the ε -noise sensitivity of the n -bit MAJORITY function as $n \rightarrow \infty$. The Majority Is Stablest Theorem was conjectured in [63] and used to prove that the approximation factor achieved by Goemans and Williamson’s algorithm for MaxCut is optimal.

Theorem 6.4 (Mossel, O’Donnell, Oleszkiewicz [83]) *Fix $0 < \varepsilon < \frac{1}{2}$. Let $f : \{-1, 1\}^n \mapsto \{-1, 1\}$ be a balanced function with all its influences at most η . then*

$$\text{NS}_\varepsilon(f) \geq \frac{1}{\pi} \arccos(1 - 2\varepsilon) - \delta,$$

where $\delta \rightarrow 0$ as $\eta \rightarrow 0$.

We present a sketch of the proof as it demonstrates the connection to an isoperimetric problem in geometry and its solution by Borell [19]. The proof involves an application of the *invariance principle*, also studied by Rotar [98] and Chatterjee [24], with Mossel’s [82] results being the state of the art. Here is a rough statement of the invariance principle:

Invariance Principle [98, 83, 24, 82]: *Suppose f is a low degree multi-linear polynomial in n variables and all its variables have small influence. Then the distribution of the values of f is nearly identical when the input is a uniform random point from $\{-1, 1\}^n$ or a random point from \mathbb{R}^n with standard Gaussian measure.*

The invariance principle allows us to translate the noise sensitivity problem on boolean hypercube to a similar problem in the Gaussian space and the latter problem has already been solved by Borell! Towards this end, let $f : \{-1, 1\}^n \mapsto \{-1, 1\}$ be a balanced boolean function whose all influences are at most η . We intend to lower bound its ε -noise sensitivity. We know that f has a representation as a multi-linear polynomial, namely its Fourier expansion:

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} x_i \quad \forall x \in \{-1, 1\}^n.$$

Let $f^* : \mathbb{R}^n \mapsto \mathbb{R}$ be a function that has the same representation as a multi-linear polynomial as f :

$$f^*(x^*) = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} x_i^* \quad \forall x^* \in \mathbb{R}^n. \quad (1)$$

Assume for the time being that f has *low* degree. Since f has all influences small, by the invariance principle, the distributions of $f(x)$ and $f^*(x^*)$ are nearly identical, and let us assume them to be exactly identical for the sake of argument. This implies that $\mathbb{E}[f^*] = \mathbb{E}[f] = 0$ and since f is boolean, so is f^* . In other words, f^* is a partition of \mathbb{R}^n into two sets of equal (Gaussian) measure. The next observation is that the ε -noise sensitivity of f is same as the ε -“Gaussian noise sensitivity” of $f^* : \mathbb{R}^n \mapsto \{-1, 1\}$. To be precise, let (x^*, y^*) be a pair of $(1 - 2\varepsilon)$ -correlated n -dimensional Gaussians, i.e. for every co-ordinate i , (x_i^*, y_i^*) are $(1 - 2\varepsilon)$ -correlated standard Gaussians. One way

to generate such a pair is to pick two independent standard n -dimensional Gaussians x^* and z^* , and let $y^* = (1 - 2\varepsilon)x^* + \sqrt{1 - (1 - 2\varepsilon)^2}z^*$, and thus one can think of y^* as a small perturbation of x^* . Let the ε -noise sensitivity of a function $f^* : \mathbb{R}^n \mapsto \{-1, 1\}$ be defined as:

$$\text{NS}_\varepsilon(f^*) := \Pr_{x^* \sim_{1-2\varepsilon} y^*} [f^*(x^*) \neq f^*(y^*)].$$

When f^* is a multi-linear polynomial as in (1), it is easily observed that

$$\text{NS}_\varepsilon(f^*) = \frac{1}{2} - \frac{1}{2} \sum_{S \subseteq [n]} \widehat{f}(S)^2 (1 - 2\varepsilon)^{|S|}.$$

But this expression is same as the ε -noise sensitivity of the boolean function f and thus $\text{NS}_\varepsilon(f) = \text{NS}_\varepsilon(f^*)$ and Theorem 6.4 follows from Borell's result that lower bounds $\text{NS}_\varepsilon(f^*)$. The parameter δ in the statement of Theorem 6.4 accounts for additive errors involved at multiple places during the argument. Also, we need to make sure that f is a low degree function in order to apply the invariance principle; this is handled by *smoothing* the function beforehand; we omit the details.

Theorem 6.5 ([19]) *If $f^* : \mathbb{R}^n \mapsto \{-1, 1\}$ is a measurable function with $\mathbb{E}[f^*] = 0$, then*

$$\text{NS}_\varepsilon(f^*) \geq \text{NS}_\varepsilon(\text{HALF SPACE}) = \frac{1}{\pi} \arccos(1 - 2\varepsilon),$$

where HALF-SPACE is the partition of \mathbb{R}^n by a hyperplane through origin.

6.5 It Ain't Over Till It's Over Theorem

For a set of co-ordinates $S \subseteq \{1, \dots, n\}$ and a string $x \in \{-1, 1\}^n$, a sub-cube $C_{S,x}$ corresponds to the set of all inputs that agree with x *outside* of S , i.e.

$$C_{S,x} := \{z \mid z \in \{-1, 1\}^n, \forall i \notin S, z_i = x_i\}.$$

A random sub-cube $C_{S,x}$ of dimension εn is picked by selecting a random set $S \subseteq \{1, \dots, n\}$, $|S| = \varepsilon n$ and a random string x . For a (roughly) balanced function f , consider the probability that f is constant on a random sub-cube of dimension εn . The dictatorship function passes the test with probability $1 - \varepsilon$ whereas it can be shown that if f has all its influences small, then f passes the test with small probability. It follows from the It Ain't Over Till It's Over Theorem of Mossel *et al* [83]. The theorem is in fact stronger: if f has all influences small, then for almost all sub-cubes C , not only that f is non-constant on C , but f takes both the values $\{-1, 1\}$ on a constant fraction of points in C . A formal statement appears below:

Theorem 6.6 ([83]) *For every $\varepsilon, \delta > 0$, there exist $\gamma, \eta > 0$ such that if $f : \{-1, 1\}^n \mapsto \{-1, 1\}$ is a $(\frac{1}{3}, \frac{2}{3})$ balanced function with all influences at most η , then*

$$\Pr_C \left[\left| \mathbb{E}[f(x) \mid x \in C] \right| \geq 1 - \gamma \right] \leq \delta,$$

where C is a random εn -dimensional sub-cube.

The theorem is proved using the invariance principle. Bansal and Khot [15] give an alternate simple proof without using the invariance principle (the random sub-cube test is proposed therein), but the conclusion is only that f is non-constant on almost every sub-cube (which suffices for their application). The random sub-cube test qualifies as a *one free bit* Long Code test, with completeness close to 1 and soundness close to zero. Via a connection between PCPs and the Independent Set problem, it implies that given an N -vertex graph with two disjoint independent sets of size $(\frac{1}{2} - \varepsilon)N$ each, it is NP-hard to find an independent set of size δN , modulo UGC [15].

6.6 Gowers Uniformity

Samorodnitsky and Trevisan [101] prove an inapproximability factor of $\Omega(2^k/k)$ for the Max- k CSP problem, which is optimal up to a multiplicative constant. Their proof relies on the analysis of the *hypergraph linearity test* (think of $k = 2^q - 1$) on a function $f : \{-1, 1\}^n \mapsto \{-1, 1\}$:

- Pick q inputs $x^1, \dots, x^q \in \{-1, 1\}^n$ at random.
- For every set $S \subseteq \{1, \dots, q\}$, $|S| \geq 2$, let $w^S := \prod_{i \in S} x^i$ and z^S be an ε -noisy version of w^S . That is, z^S is obtained by flipping every bit of w^S with probability ε independently.
- Accept if and only if for every S , $f(z) = \prod_{i \in S} f(x^i)$.

If f is a dictatorship, it is easily seen that it passes the test with probability at least $1 - \varepsilon \cdot 2^q$. This is because it would pass the test with probability 1 if there were no noise, there are $2^q - q - 1$ sets S , and the test for each S could fail with probability ε due to the noise.

On the other hand, if f is a balanced function and all its influences are at most η , it can be shown that f passes the test with probability at most $\frac{1}{2^{2^q - q - 1}} + \delta$ where $\delta \rightarrow 0$ as $\eta \rightarrow 0$. Note that there are $2^q - q - 1$ sub-tests, one for each $|S| \geq 2$. If f has all influences small, these tests behave as if they were independent tests, each test accepts with probability essentially $\frac{1}{2}$, and hence the probability that all tests accept is close to $\frac{1}{2^{2^q - q - 1}}$.

Samorodnitsky and Trevisan [101] relate the acceptance probability of the test to the Gowers Uniformity norms [45] of a function, and then show that for a function with all influences small, the Gowers Uniformity norm is small as well.

Definition 6.7 Gowers Uniformity: *Let $f : \{-1, 1\}^n \mapsto \{-1, 1\}$ be a function, and $d \geq 1$ be an integer. The dimension- d uniformity of f is defined as:*

$$U^d(f) := \mathbb{E}_{x, x^1, \dots, x^d} \left[\prod_{S \subseteq \{1, \dots, d\}} f \left(x \cdot \prod_{i \in S} x^i \right) \right].$$

Theorem 6.8 ([101]) *If f is a balanced function and $\forall i \in [n]$, $\text{Inf}_i(f) \leq \eta$, then $U^d(f) \leq \sqrt{\eta} \cdot 2^{O(d)}$.*

6.7 Some Open Problems

We conclude this section by listing some open Fourier analytic problems that arise from applications to inapproximability.

- The MaxCut problem can be generalized to the Max- k Cut problem where one seeks to partition a graph into $k \geq 3$ sets so as to maximize the number of edges cut. An optimal inapproximability result for this problem is implied by the Plurality Is Stablest Conjecture stating that the Plurality function from $\{1, \dots, k\}^n$ to $\{1, \dots, k\}$ is the most stable under noise among the class of functions that are balanced and whose all influences are small.¹⁰ Using the invariance principle, this conjecture can be reduced to the **Standard Simplex Conjecture** stated in Section 8.1 (see [63, 53] for a detailed exposition where the conjectures are made).

¹⁰There is a natural generalization for the notions of influence and Fourier representation of k -ary functions on a k -ary hypercube that we don't state here.

- Khot and Naor [64, 65] study the Kernel Clustering problem and the following question arises from their work: if a function $f : \{1, \dots, k\}^n \mapsto \{1, \dots, k\}$ has all influences small, what is the maximum Fourier mass the function can have at the first level? Again using the invariance principle, the question can be reduced to a geometric problem described in Section 8.1. The question is easily answered for $k = 2, 3$, but remains open for $k \geq 4$.
- Consider a variant of the test discussed in Section 6.5: Assume that f is balanced, and one tests whether f is constant -1 on a random sub-cube of linear dimension. We know that if f passes the test with constant probability, say α , then it must have an influential variable. However f need not be close to a junta.¹¹ Is it necessarily true that there is a function g that is close to a junta, monotonically above f , and passes the test with probability close to α ? We say that g is monotonically above f if $\forall x, f(x) = 1 \implies g(x) = 1$. Such a result, though interesting on its own, might be useful towards inapproximability of problems related to graph coloring.

7 Parallel Repetition Theorem

Let $\mathcal{U}_{2p1r}(G(V, W, E), [m], [n], \{\pi_e | e \in E\})$ be an instance of the 2-Prover-1-Round Game problem as in Definition 2.2. For an integer $k \geq 1$, the k -wise parallel repetition game is defined as the instance $\mathcal{U}_{2p1r}^{\otimes k}(G'(V', W', E'), [m'], [n'], \{\pi_{e'} | e' \in E'\})$ where,

- $V' := V^k$, $W' := W^k$. For $v' = (v_1, \dots, v_k) \in V^k$, $w' = (w_1, \dots, w_k) \in W^k$, we have $e' = (v', w') \in E'$ if and only if $e_i = (v_i, w_i) \in E \forall i \in [k]$.
- $[m']$ is identified with $[m]^k$ and $[n']$ is identified with $[n]^k$.
- If $e' = (v', w') \in E'$, then $\pi_{e'} : [m'] \mapsto [n']$ is defined as:

$$\pi_{e'}(i_1, \dots, i_k) := (\pi_{e_1}(i_1), \dots, \pi_{e_k}(i_k)).$$

It is clear that $\text{OPT}(\mathcal{U}_{2p1r}^{\otimes k}) \geq \text{OPT}(\mathcal{U}_{2p1r})^k$ by considering a labeling to the repeated game that is, on every co-ordinate, same as the optimal labeling to the original game. The parallel repetition theorems give an upper bound on the optimum of the repeated game in terms of that of the original game.

Theorem 7.1 (Raz [96]) *If $\text{OPT}(\mathcal{U}_{2p1r}) = 1 - \varepsilon$, then*

$$\text{OPT}(\mathcal{U}_{2p1r}^{\otimes k}) \leq (1 - \Omega(\varepsilon^c))^{\frac{k}{\log m + \log n}},$$

where m and n are the sizes of the label sets and c is an absolute constant.

It is implicit in Raz's paper that the exponent $c = 32$ works. Holenstein [52] simplified Raz's proof and also improved the exponent to $c = 3$. Both the results apply to more general 2-Prover-1-Round games, and not just the projection games that are focus of this article. For the case of projection games, Rao [95] showed that the exponent can be improved to $c = 2$ and moreover there is no dependence on the size of the label sets.

¹¹ A junta is a function that depends on a bounded number of co-ordinates.

Theorem 7.2 (Rao [95]) *If $\text{OPT}(\mathcal{U}_{2p1r}) = 1 - \varepsilon$, then*

$$\text{OPT}(\mathcal{U}_{2p1r}^{\otimes k}) \leq (1 - \Omega(\varepsilon^2))^k.$$

As mentioned before, Rao’s Theorem can be used to amplify the gap of a Unique Game instance and show that the Unique Games Conjecture is implied by its weak version.

Theorem 7.3 ([95]) *Conjecture 3.1 implies Conjecture 2.5.*

Proof: Assume Conjecture 3.1 and \mathcal{U} be the Unique Game instance therein. Let $k := \frac{1}{\varepsilon \cdot \Gamma(1/\varepsilon)}$ and consider the instance $\mathcal{U}^{\otimes k}$. Note that the parallel repetition of a unique game is also a unique game.¹²

- (YES case:) $\text{OPT}(\mathcal{U}) \geq 1 - \varepsilon$, and hence $\text{OPT}(\mathcal{U}^{\otimes k}) \geq (1 - \varepsilon)^k \geq 1 - k \cdot \varepsilon = 1 - \frac{1}{\Gamma(1/\varepsilon)}$.
- (NO Case:) $\text{OPT}(\mathcal{U}) \leq 1 - \Omega(\sqrt{\varepsilon} \cdot \Gamma(1/\varepsilon))$, and by Theorem 7.2,

$$\text{OPT}(\mathcal{U}^{\otimes k}) \leq (1 - \Omega(\varepsilon \cdot \Gamma(1/\varepsilon)^2))^{\frac{1}{\varepsilon \cdot \Gamma(1/\varepsilon)}} \leq \exp(-\Omega(\Gamma(1/\varepsilon))).$$

Thus as $\varepsilon \rightarrow 0$, the optimum of the repeated game tends to 1 in the YES Case and to 0 in the NO Case. ■

Arguably, stronger parallel repetition theorems should be useful towards gap amplification of unique games. One question is whether in Theorem 7.2, the exponent 2 can be further improved to $c < 2$. If this were the case, one would have the equivalence between the UGC and an appropriate inapproximability result for the MaxCut problem.

Theorem 7.4 ([63]) *The UGC implies that for all sufficiently small $\varepsilon > 0$, $\text{GapMaxCut}_{1-\varepsilon, 1-b\sqrt{\varepsilon}}$ is NP-hard where $b = \frac{2}{\pi}$.*

Now suppose Theorem 7.2 were correct with the exponent $c < 2$. Then, starting with an instance of $\text{GapMaxCut}_{1-\varepsilon, 1-b\sqrt{\varepsilon}}$ (which is a unique game), one could amplify the gap using such a strong parallel repetition theorem and obtain an instance of $\text{GapUG}_{1-\varepsilon', \delta'}$. This would show that the UGC is equivalent to saying that $\text{GapMaxCut}_{1-\varepsilon, 1-b\sqrt{\varepsilon}}$ is NP-hard. However, Raz [97] showed that the exponent $c = 2$ in Theorem 7.2 cannot be improved further. The *odd cycle game* serves as a counter-example.

The odd cycle game is a two-prover game, first suggested in [32] and further motivated and studied in [38, 14]. Let $\ell \geq 3$ be an odd integer and consider a graph that is a single cycle of length ℓ . Intuitively, the two provers are trying to convince the verifier that the graph is 2-colorable. With probability one half the verifier asks the two provers about the color of the same node in the graph and accepts if their answers are the same. With probability one half the verifier asks the two provers about the colors of two adjacent nodes in the graph and accepts if their answers are different. It is easy to show that the optimum (value) of the odd cycle game is $1 - \frac{1}{2\ell}$. Feige, Kindler, and O’Donnell [38] studied the parallel repetition of this game and showed that the value of the k -wise repeated game is at most $1 - \tilde{\Omega}\left(\frac{\sqrt{k}}{\ell}\right)$. Raz [97] showed that this is essentially the tight bound.

¹²Strictly speaking, the parallel repetition theorem applies to 2-Prover-1-Round games (which are “bipartite”), whereas we defined Unique Game problem as a constraint satisfaction problem on a general graph. This distinction is minor in our context, and if one desires, one can switch back and forth between the two formulations as indicated in the paragraph after Definition 2.4.

Theorem 7.5 (Raz [97]) *For $k \leq \ell^2$, there is a prover strategy for the odd cycle game (on an ℓ -cycle) that makes the verifier accept with probability $1 - O\left(\frac{\sqrt{k}}{\ell}\right)$.*

The theorem implies that for $k \geq \ell^2$, the value of the k -wise repeated game is at least $(1 - \frac{1}{\ell^2})^{O(k)}$ showing that the exponent $c = 2$ in Theorem 7.2 cannot be further improved. Feige *et al* [38] also discovered a beautiful connection between the value of the repeated odd cycle game and a tiling of Euclidean space, see Section 8.3.

Feige and Lovasz [39] observed a connection between the SDP relaxation of a Unique Game instance, e.g. the relaxation in Figure 4 - Section 5, and its behavior under parallel repetition.¹³ Consider a Unique Game instance $\mathcal{U}(G(V, E), [n], \{\pi_e | e \in E\})$ and let $\text{OPT}(\mathcal{U})$ and $\text{SDP}(\mathcal{U})$ denote the optimal value of the game and the SDP relaxation respectively. Feige and Lovasz showed that the SDP optimum is multiplicative under parallel repetition:

Theorem 7.6 ([39]) $\text{SDP}(\mathcal{U}^{\otimes k}) = \text{SDP}(\mathcal{U})^k$.

Note that this theorem implies a weak parallel repetition theorem for unique games: Suppose a Unique Game has value bounded away from 1. It easily implies that the SDP optimum is also bounded away from 1. Since the SDP value is multiplicative and is an upper bound on the value of a game, it follows that the value of the k -wise repeated game tends to zero as $k \rightarrow \infty$.

Barak *et al* [17] used techniques from Raz’s paper to construct prover strategies and give a lower bound on the value of a repeated game in terms of the SDP optimum of the original game:

Theorem 7.7 ([17]) *If $\text{SDP}(\mathcal{U}) \geq 1 - \beta$, then:*

- $\text{OPT}(\mathcal{U}^{\otimes k}) \geq 1 - O(\sqrt{k\beta \log(n/\beta)})$.
- $\text{OPT}(\mathcal{U}^{\otimes k}) \geq 1 - O(\sqrt{k\beta})$ for unique games with $n = 2$ labels (e.g. MaxCut).

Consider the following approach towards proving UGC: first prove hardness of $\text{GapMaxCut}_{1-\varepsilon, 1-b\sqrt{\varepsilon}}$ and then amplify hardness by parallel repetition. As we discussed before, the exponent 2 in Theorem 7.2 cannot be improved, and therefore the amplification step is stuck. Theorem 7.7 in fact gives an inherent (and subtle) reason why this approach is doomed to fail. Suppose on the contrary that, we first showed $\text{GapMaxCut}_{1-\varepsilon, 1-b\sqrt{\varepsilon}}$ to be NP-hard and then amplified the gap, via k -wise parallel repetition, to $\text{GapUG}_{1-\varepsilon', \delta'}$. In the YES Case, we would argue that the optimum of the repeated instance is at least $(1 - \varepsilon)^k$, and since this is required to be close to 1, we must have $\varepsilon k \ll 1$. On the other hand, since $\text{GapMaxCut}_{1-\varepsilon, 1-b\sqrt{\varepsilon}}$ is NP-hard, an SDP relaxation would not be able to distinguish between the gap, and hence there must exist MaxCut instances G such that $\text{OPT}(G) \leq 1 - b\sqrt{\varepsilon}$ (i.e. the NO Case) but $\text{SDP}(G) \geq 1 - \varepsilon$. Theorem 7.7 then implies that $\text{OPT}(G^{\otimes k}) \geq 1 - O(\sqrt{k\varepsilon}) \approx 1$. Thus, there are NO instances such that the repeated game has value close to 1 rather than close to 0!

Note that we currently have no approach to prove even the “starting” result, i.e. the hardness of $\text{GapMaxCut}_{1-\varepsilon, 1-b\sqrt{\varepsilon}}$. In author’s opinion, such a result would be a huge step towards proving the UGC (especially in light of the equivalence between UGC and Conjecture 3.1), even though it might not directly imply the UGC as far as the current state of knowledge is given.

¹³ The Feige-Lovasz relaxation is somewhat different and the parallel repetition typically refers to only bipartite games. We ignore these subtleties here.

8 Geometry

The connections between the Unique Games Conjecture and geometry arise in (at least) three contexts: proofs of Fourier analytic theorems via the use of the invariance principle, constructing integrality gap instances, and rather unexpectedly, constructing “foams”. We explain these connections in three separate sections.

8.1 Proofs of Fourier Analytic Theorems

Raghavendra’s Theorem 4.2 shows that for every CSP, the integrality gap of a natural SDP-relaxation gives the optimal approximation threshold for that CSP (modulo UGC). Characterizing the integrality gap is essentially a geometric question and we point out a couple of examples where the geometric question is interesting in its own right (in author’s opinion).

The Standard Simplex Conjecture: As demonstrated in Section 6.4, the MaxCut problem motivates the Majority Is Stablest Theorem which in turn, reduces to an isoperimetric question in Gaussian space, namely the question of partitioning \mathbb{R}^n into two sets of equal (Gaussian) measure so as to minimize the Gaussian noise-sensitivity. The latter question has been answered by Borell’s Theorem characterizing a half-space through the origin as an optimal partition.

We also mentioned the Max- k Cut problem and the *Plurality is Stablest Conjecture* that it motivates. Using the invariance principle, this Fourier analytic conjecture can be reduced to the following geometric conjecture (see Isaksson and Mossel [53]).

The Standard Simplex Conjecture: *The standard k -simplex partition is the most noise-stable balanced partition of \mathbb{R}^n with $n \geq k - 1$.*

A partition of \mathbb{R}^n into k measurable sets A_1, \dots, A_k is called balanced if each A_i has (Gaussian) measure $\frac{1}{k}$. The ε -noise sensitivity is defined as the probability that two $(1 - 2\varepsilon)$ -correlated n -dimensional standard Gaussian points $x, y \in \mathbb{R}^n$ belong to different sets in the partition. The standard k -simplex partition of \mathbb{R}^n is obtained by letting $\mathbb{R}^n = \mathbb{R}^{k-1} \times \mathbb{R}^{n-k+1}$ and then partitioning \mathbb{R}^{k-1} into k regular simplicial cones.

The Propeller Conjecture: In Section 6.7, we mentioned the Kernel Clustering problem and the corresponding Fourier question of maximizing the Fourier mass at the first level among the class of low-influence functions. Using invariance principle, the Fourier question can be reduced to the following geometric question:

Definition 8.1 *Let A_1, \dots, A_k be a partition of \mathbb{R}^{k-1} into k measurable sets and for $1 \leq \ell \leq k$, let z_ℓ be the Gaussian moment vector over A_ℓ , i.e.*

$$z_\ell := \int_{A_\ell} x \, d\gamma \quad \text{where } \gamma \text{ is standard Gaussian measure on } \mathbb{R}^{k-1}.$$

Let $C(k)$ be the supremum (it is achieved) of the sum of squared lengths of z_ℓ ’s over all possible partitions, i.e.

$$C(k) := \sup_{\mathbb{R}^{k-1} = A_1 \cup \dots \cup A_k} \sum_{\ell=1}^k \|z_\ell\|^2. \tag{2}$$

Note that there is no restriction on the measures of the k sets in the partition and in particular, the partition need not be balanced. It seems challenging to characterize an optimal partition (in the sense of (2)) for $k \geq 4$. For $k = 2$, the optimal partition of \mathbb{R} into two sets is the partition into

positive and negative real line, and $C(2) = \frac{1}{\pi}$. For $k = 3$, the optimal partition of \mathbb{R}^2 into three sets is the “propeller”, i.e. partition into three cones with angle 120° each, and $C(3) = \frac{9}{8\pi}$. One would expect that for $k = 4$, the optimal partition of \mathbb{R}^3 into four sets is the partition into four cones given by a regular tetrahedron. This turns out to be false as numerical computation shows that the value of this partition is worse than $C(3) = \frac{9}{8\pi}$ that can be achieved by letting $\mathbb{R}^3 = \mathbb{R}^2 \times \mathbb{R}$ and then partitioning \mathbb{R}^2 as a propeller. In fact Khot and Naor [64, 65] conjecture that the propeller partition is the optimal one for any $k \geq 3$:

Conjecture 8.2 Propeller Conjecture: *For every $k \geq 3$, $C(k) = C(3)$. In words, the optimal partition of \mathbb{R}^{k-1} into k sets in the sense of (2) is achieved by letting $\mathbb{R}^{k-1} = \mathbb{R}^2 \times \mathbb{R}^{k-3}$ and partitioning \mathbb{R}^2 into a propeller.*

8.2 Integrality Gaps

It was demonstrated in [71] that the Unique Games Conjecture leads to explicit constructions of integrality gap instances. We briefly describe this approach. Suppose we have a gap-preserving reduction from the Unique Game problem to a maximization problem \mathcal{I} such as MaxCut. Specifically, the reduction maps a Unique Game instance \mathcal{U} to an instance I such that for some $0 < s < c$,

- (YES Case:) $\text{OPT}(\mathcal{U}) \geq 1 - \varepsilon \implies \text{OPT}(I) \geq c$.
- (NO Case:) $\text{OPT}(\mathcal{U}) \leq \delta \implies \text{OPT}(I) \leq s$.

Such a reduction can then be used to map an integrality gap instance for the Unique Game problem to that of \mathcal{I} in the following way. We construct an instance \mathcal{U} of the Unique Game problem such that w.r.t. the SDP-relaxation in Figure 4, Section 5,

$$\text{SDP}(\mathcal{U}) \geq 1 - \varepsilon \quad \text{and} \quad \text{OPT}(\mathcal{U}) \leq \delta. \tag{3}$$

Now we use the reduction to map \mathcal{U} to an instance I . The reduction guarantees that $\text{OPT}(I) \leq s$ (the NO Case). On the other hand, a close examination of the reduction shows that there is a way to map the SDP solution for \mathcal{U} to a corresponding SDP solution for I that “mimics” the YES Case, i.e. if the SDP solution for \mathcal{U} has value $1 - \varepsilon$, then the SDP solution for I that we construct has value at least c . This gives us the desired integrality gap instance for \mathcal{I} . This approach is quite general and applies to every CSP. The MaxCut and Sparsest Cut are perhaps the most notable illustrations as described below.

The construction of the Unique Game integrality gap instance as in (3) [71] is rather interesting. We don’t describe it here, but do point out that it is hypercube based and the fact that the (integral) optimum is low is related to the fact that small subsets of a boolean hypercube are almost fully expanding. This partly motivates Theorem 5.2 giving a connection between the UGC and the Small Set Expansion problem [94].

MaxCut with Triangle Inequalities: As mentioned in Section 2.6, Feige and Schechtman [41] already constructed integrality gap instances for the Goemans-Williamson’s MaxCut SDP, with gap arbitrarily close to α_{MC} . As shown in [63], the UGC implies that it is NP-hard to approximate MaxCut within any factor strictly less than α_{MC} . Therefore, modulo the UGC (and $\text{P} \neq \text{NP}$), even if one added polynomially many constraints to the SDP in Figure 1, the integrality gap would still remain as α_{MC} (as otherwise one would have a better polynomial time algorithm). However, this argument is conditional on UGC, and it was an open problem whether addition of the following *triangle inequality constraints* to the SDP in Figure 1 improves the MaxCut integrality gap:

$$\forall i, j, k, \in [N], \quad \|v_i - v_j\|^2 + \|v_j - v_k\|^2 \geq \|v_i - v_k\|^2. \tag{4}$$

Note that if v_i, v_j, v_k are $\{-1, 1\}$ valued integer variables, then the constraints are indeed satisfied and therefore these are legal constraints in an SDP-relaxation. Using the Unique Game based strategy above, Khot and Vishnoi [71] were able to construct integrality gap instances for Max-Cut with gap arbitrarily close to α_{MC} and where the SDP vectors also obey the triangle inequality constraints.

Sparsest Cut and Connection to Metric Embeddings: As shown in [25, 71], the UGC implies that it is NP-hard to approximate the (non-uniform) Sparsest Cut problem within any constant factor. On the other hand, one can write a natural SDP relaxation for the problem, very similar to Figure 1, where one intends to minimize the fraction of edges cut and adds a constraint that the cut is *balanced with respect to demands*. The triangle inequality constraints as in (4) are added as well. Characterizing the integrality gap for this SDP has been a challenging open problem for many years. It turns out that the integrality gap on N -vertex graphs is exactly the least number $c_1(\text{NEG}, N)$ such that every N -point “negative type” metric embeds into the class of ℓ_1 metrics with *distortion* $c_1(\text{NEG}, N)$, i.e. while preserving all distances up to a factor of $c_1(\text{NEG}, N)$. This follows from the characterization of the ℓ_1 metrics as convex combinations of “cut-metrics” and from the fact that the SDP solution yields a negative-type metric thanks to the triangle inequality constraints. A metric $d(\cdot, \cdot)$ is called negative-type if its square root $\sqrt{d(\cdot, \cdot)}$ (which is also a metric) isometrically embeds into ℓ_2 . These are precisely the metrics that arise as $d(i, j) = \|v_i - v_j\|^2$ for vectors $\{v_i\}_{i=1}^N$ satisfying (4).

Goemans and Linial [44, 79] conjectured that $c_1(\text{NEG}, N)$ is a universal constant independent of N . This would yield a constant factor approximation to the Sparsest Cut problem, in contradiction with the UGC-based inapproximability result [25, 71]. In fact, using the strategy described before, Khot and Vishnoi [71] were able to construct an integrality gap instance for the Sparsest Cut problem with a super-constant gap and disprove the Goemans and Linial conjecture:

Theorem 8.3 ([71]) $c_1(\text{NEG}, N) \geq \Omega((\log \log N)^c)$ for some constant $c > 0$.

The lower bound is strengthened by Raghavendra and Steurer [92] (also by Khot and Saket [70] with a quantitatively weaker result):

Theorem 8.4 ([92, 70]) *There is an N -point negative type metric such that its submetric on any subset of t points is isometrically ℓ_1 -embeddable, but the whole metric incurs distortion of at least t to embed into ℓ_1 , and $t = (\log \log N)^c$ for some constant $c > 0$.*

The [92, 70] constructions again use the Unique Game-based strategy, in even more sophisticated way. They first construct an integrality gap instance for the Unique Game SDP relaxation augmented with t rounds of the so-called Sherali-Adams linear programming hierarchy. The instance is then translated to an integrality gap instance for the Sparsest Cut SDP augmented with t rounds of Sherali-Adams. The latter result is same as a local versus global ℓ_1 embeddability result for the class of negative type metrics.

We also mention related works that are not necessarily related to the UGC. In a breakthrough work, Arora, Rao, and Vazirani [10] showed that the SDP algorithm gives $O(\sqrt{\log N})$ approximation to the Sparsest Cut problem without demands. This was extended to the demands version of the problem by Arora, Lee, and Naor [8], albeit with a slight loss in the approximation factor. As discussed, the result is equivalent to an upper bound on $c_1(\text{NEG}, N)$.

Theorem 8.5 ([8]) $c_1(\text{NEG}, N) \leq \tilde{O}(\sqrt{\log N})$.

Using an alternate construction based on the geometry of Heisenberg group, a sequence of works by Lee and Naor [76], Cheeger and Kleiner [26, 27], Cheeger, Kleiner, and Naor [28, 29] obtained a stronger lower bound than Theorem 8.3 (see Naor’s survey for an excellent survey on this topic [85]):

Theorem 8.6 ([76, 26, 27, 28, 29]) $c_1(\text{NEG}, N) \geq \Omega((\log N)^c)$ for some constant $c > 0$.

8.3 Constructing a Foam

Recall the odd cycle game from Section 7. Feige *et al* [38] showed that the question of characterizing the value of a repeated odd cycle game is related to the question of constructing a “foam”, i.e. a tiling of space using a shape/cell of unit volume while minimizing its surface area!

A Foam Problem: Let $S \subseteq \mathbb{R}^d$ be a “shape” with unit volume such that its translations by \mathbb{Z}^d tile \mathbb{R}^d . What is the least surface area of such a shape?

Clearly, the surface area of a tiling shape S must at least be that of the unit volume ball, namely $\Omega(\sqrt{d})$. But the unit volume ball is not a legal tiling shape. On the other hand, a unit volume cube is a legal tiling shape, and has surface area of d . The best known constructions were only a constant factor better than the cube. Kindler, O’Donnell, Rao, and Wigderson [73] used a continuous version of Raz’s counter-example, Theorem 7.5, to construct a tiling shape whose surface area is $O(\sqrt{d})!$ Alon and Klartag [2] subsequently gave a more direct construction.

9 Conclusion

We conclude by citing arguments in favor and against the Unique Games Conjecture. None of these are particularly compelling in the author’s opinion and are subject to personal bias.

Arguments in Favor:

- The conjecture states that $\text{GapUG}_{1-\varepsilon, \delta}$ is NP-hard. Feige and Reichman [40] show that for every constant $\delta > 0$, there is a constant $C(\delta) > 1$ such that $\text{GapUG}_{C(\delta) \cdot \delta, \delta}$ is NP-hard, and $C(\delta) \rightarrow \infty$ as $\delta \rightarrow 0$. However, in their result, $C(\delta) \cdot \delta \rightarrow 0$ as $\delta \rightarrow 0$ and the result falls short of proving the UGC.
- Conjecture 3.1, which is a priori much weaker than the UGC, seems much more believable, but is in fact equivalent to the UGC.
- The SDP relaxation (Figure 4) for the Unique Game problem has an $(1 - \varepsilon, \delta)$ integrality gap.
- The UGC implies a super-constant inapproximability result for the Sparsest Cut problem and can be used to explicitly construct super-constant integrality gap for its SDP relaxation, contrary to the Goemans-Linial Conjecture.

Arguments Against:

- We do not know any natural distribution on Unique Game instances such that the problem is hard for the known algorithms on this distribution.
- We do not know any other natural problem that is provably equivalent to the Unique Game problem.

- We do not know Unique Game integrality gap instances for the so-called Lasserre SDP relaxation. The current state of knowledge does not rule out the possibility that a Lasserre SDP, even with a small (say three) number of rounds, may disprove the UGC. See however [67] for a construction of an *approximate* Lasserre integrality gap.

As wishful thinking, it would be “aesthetically pleasing” to have the UGC proven correct, have all the “right” inapproximability results in a unifying way, and all the connections to analysis and geometry nicely fit together. On the other hand, proving the conjecture incorrect would be an algorithmic breakthrough, likely going beyond the SDP barrier, since the UGC seems to exactly capture the limitations of SDPs.¹⁴ Irrespective of its truth, the conjecture will probably lead to even more techniques and unconditional results.

References

- [1] A. Agarwal, M. Charikar, K. Makarychev, and Y. Makarychev. $O(\sqrt{\log n})$ approximation algorithms for min-UnCut, min-2CNF deletion, and directed cut problems. In *Proc. 37th ACM Symposium on Theory of Computing*, pages 573–581, 2005.
- [2] N. Alon and B. Klartag. Economical toric spines via Cheeger’s inequality. *Journal of Topology and Analysis*, 1:101–111, 2009.
- [3] C. Ambühl, M. Mastrolilli, and O. Svensson. Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling. In *Proc. Annual IEEE Symposium on Foundations of Computer Science*, pages 329–337, 2007.
- [4] S. Arora, L. Babai, J. Stern, and E. Sweedyk. The hardness of approximate optima in lattices, codes and systems of linear equations. *Journal of Computer and Systems Sciences*, 54:317–331, 1997.
- [5] S. Arora, B. Barak, and D. Steurer. A $2^{n^{\text{poly}(\epsilon)}}$ -time algorithm for $(1 - \epsilon)$ -satisfiable unique games. *Manuscript*, 2010.
- [6] S. Arora, E. Chlamtac, and M. Charikar. New approximation guarantee for chromatic number. In *Proc. ACM Symposium on the Theory of Computing*, pages 215–224, 2006.
- [7] S. Arora, S. Khot, A. Kolla, D. Steurer, M. Tulsiani, and N. K. Vishnoi. Unique games on expanding constraint graphs are easy. In *Proc. ACM Symposium on the Theory of Computing*, pages 21–28, 2008.
- [8] S. Arora, J. Lee, and A. Naor. Euclidean distortion and the sparsest cut. In *Proc. 37th ACM Symposium on Theory of Computing*, pages 553–562, 2005.
- [9] S. Arora, C. Lund, R. Motawani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.
- [10] S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings and graph partitioning. In *Proc. 36th ACM Symposium on Theory of Computing*, pages 222–231, 2004.

¹⁴It may be the case that the $\text{GapUG}_{1-\epsilon, \delta}$ problem is “hard”, but not necessarily NP-hard. In this case, in author’s opinion, the UGC would still be “morally correct”, especially since its main motivation is to prove inapproximability results.

- [11] S. Arora and S. Safra. Probabilistic checking of proofs : A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.
- [12] P. Austrin. Balanced max 2-sat might not be the hardest. In *Proc. ACM Symposium on the Theory of Computing*, pages 189–197, 2007.
- [13] P. Austrin and E. Mossel. Approximation resistant predicates from pairwise independence. *Computational Complexity*, 18(2):249–271, 2009.
- [14] K. Azimian and M. Szegedy. Parallel repetition of the odd cycle game. In *LATIN*, pages 676–686, 2008.
- [15] N. Bansal and S. Khot. Optimal long code test with one free bit. In *Proc. 50th IEEE Symposium on Foundations of Computer Science*, 2009.
- [16] N. Bansal and S. Khot. Inapproximability of hypergraph vertex cover and applications to scheduling problems. *Manuscript*, 2010.
- [17] B. Barak, M. Hardt, I. Haviv, A. Rao, O. Regev, and D. Steurer. Rounding parallel repetitions of unique games. In *Proc. Annual IEEE Symposium on Foundations of Computer Science*, pages 374–383, 2008.
- [18] M. Bellare, O. Goldreich, and M. Sudan. Free bits, PCPs and non-approximability. *Electronic Colloquium on Computational Complexity, Technical Report TR95-024*, 1995.
- [19] C. Borell. Geometric bounds on the Ornstein-Uhlenbeck velocity process. *Z. Wahrsch. Verw. Gebiete*, 70(1):1–13, 1985.
- [20] J. Bourgain. On the distribution of the Fourier spectrum of boolean functions. *Israel J. of Math.*, (131):269–276, 2002.
- [21] G. Calinescu, H. J. Karloff, and Y. Rabani. An improved approximation algorithm for multiway cut. *J. Comput. Syst. Sci.*, 60(3):564–574, 2000.
- [22] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for unique games. In *Proc. ACM Symposium on the Theory of Computing*, pages 205–214, 2006.
- [23] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for maximum constraint satisfaction problems. In *SODA*, 2007.
- [24] S. Chatterjee. A simple invariance theorem. *arXiv:math/0508213v1*, 2005.
- [25] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. In *Proc. 20th IEEE Conference on Computational Complexity*, pages 144–153, 2005.
- [26] J. Cheeger and B. Kleiner. Differentiating maps into L^1 and the geometry of BV functions. *To appear in Ann. Math., preprint available at <http://arxiv.org/abs/math/0611954>*, 2006.
- [27] J. Cheeger and B. Kleiner. On the differentiation of Lipschitz maps from metric measure spaces to Banach spaces. *Inspired by S.S. Chern, Volume 11 of Nankai Tracts. Math.*, pages 129–152, 2006.

- [28] J. Cheeger, B. Kleiner, and A. Naor. Compression bounds for Lipschitz maps from the Heisenberg group to l_1 . *Preprint*, 2009.
- [29] J. Cheeger, B. Kleiner, and A. Naor. A $(\log n)^{\Omega(1)}$ integrality gap for the sparsest cut sdP. In *Proc. 50th IEEE Symposium on Foundations of Computer Science*, 2009.
- [30] E. Chlamtac, K. Makarychev, and Y. Makarychev. How to play unique games using embeddings. In *Proc. Annual IEEE Symposium on Foundations of Computer Science*, pages 687–696, 2006.
- [31] J. Chuzhoy and S. Khanna. Polynomial flow-cut gaps and hardness of directed cut problems. *Journal of the ACM*, 56(2), 2009.
- [32] R. Cleve, P. Høyer, B. Toner, and J. Watrous. Consequences and limits of nonlocal strategies. In *IEEE Conference on Computational Complexity*, pages 236–249, 2004.
- [33] N. Devanur, S. Khot, R. Saket, and N. Vishnoi. Integrality gaps for sparsest cut and minimum linear arrangement problems. In *Proc. 38th ACM Symposium on Theory of Computing*, 2006.
- [34] I. Dinur, V. Guruswami, S. Khot, and O. Regev. A new multilayered PCP and the hardness of hypergraph vertex cover. In *Proc. 34th ACM Symposium on Theory of Computing*, 2002.
- [35] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. *SIAM J. Comput.*, 39(3):843–873, 2009.
- [36] I. Dinur and S. Safra. The importance of being biased. In *Proc. 34th Annual ACM Symposium on Theory of Computing*, 2002.
- [37] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Interactive proofs and the hardness of approximating cliques. *Journal of the ACM*, 43(2):268–292, 1996.
- [38] U. Feige, G. Kindler, and R. O’Donnell. Understanding parallel repetition requires understanding foams. In *Proc. Annual IEEE Conference on Computational Complexity*, pages 179–192, 2007.
- [39] U. Feige and L. Lovász. Two-prover one-round proof systems, their power and their problems. In *Proc. 24th Annual ACM Symposium on Theory of Computing*, pages 733–744, 1992.
- [40] U. Feige and D. Reichman. On systems of linear equations with two variables per equation. In *APPROX-RANDOM*, pages 117–127, 2004.
- [41] U. Feige and G. Schechtman. On the integrality ratio of semidefinite relaxations of max cut. In *Proc. ACM Symposium on the Theory of Computing*, pages 433–442, 2001.
- [42] E. Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):27–35, 1998.
- [43] M. Goemans and D. Williamson. 0.878 approximation algorithms for MAX-CUT and MAX-2SAT. In *Proc. 26th ACM Symposium on Theory of Computing*, pages 422–431, 1994.
- [44] M. X. Goemans. Semidefinite programming in combinatorial optimization. *Math. Program.*, 79:143–161, 1997.

- [45] T. Gowers. A new proof of Szemerédi’s theorem for progressions of length four. *Geometric and Functional Analysis*, 8(3):529–551, 1998.
- [46] A. Gupta and K. Talwar. Approximating unique games. In *Proc. ACM-SIAM Symposium on Discrete Algorithms*, pages 99–106, 2006.
- [47] V. Guruswami, J. Hastad, and M. Sudan. Hardness of approximate hypergraph coloring. In *Proc. 41st IEEE Symposium on Foundations of Computer Science*, pages 149–158, 2000.
- [48] V. Guruswami, R. Manokaran, and P. Raghavendra. Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In *Proc. Annual IEEE Symposium on Foundations of Computer Science*, pages 573–582, 2008.
- [49] J. Hastad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.
- [50] J. Hastad. On a protocol possibly useful for MIN-2SAT. *Manuscript*, 2001.
- [51] J. Hastad. Some optimal inapproximability results. *Journal of ACM*, 48:798–859, 2001.
- [52] T. Holenstein. Parallel repetition: simplifications and the no-signaling case. In *Proc. ACM Symposium on the Theory of Computing*, pages 411–419, 2007.
- [53] M. Isaksson and E. Mossel. New maximally stable gaussian partitions with discrete applications. *arXiv:0903.3362v3*, 2010.
- [54] J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *Proc. 29th Symposium on the Foundations of Computer Science*, pages 68–80, 1988.
- [55] D. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. In *Proc. 35th IEEE Symposium on Foundations of Computer Science*, pages 2–13, 1994.
- [56] D. R. Karger, P. N. Klein, C. Stein, M. Thorup, and N. E. Young. Rounding algorithms for a geometric embedding of minimum multiway cut. *Math. Oper. Res.*, 29(3):436–461, 2004.
- [57] H. J. Karloff. How good is the Goemans-Williamson max cut algorithm? In *Proc. ACM Symposium on the Theory of Computing*, pages 427–434, 1996.
- [58] S. Khanna, N. Linial, and S. Safra. On the hardness of approximating the chromatic number. In *Proc. 2nd Israel Symposium on Theory and Computing Systems, ISTCS*, pages 250–260, 1993.
- [59] S. Khot. Improved inapproximability results for maxclique, chromatic number and approximate graph coloring. In *Proc. 42nd IEEE Annual Symposium on Foundations of Computer Science*, 2001.
- [60] S. Khot. On the power of unique 2-prover 1-round games. In *Proc. 34th ACM Symposium on Theory of Computing*, 2002.
- [61] S. Khot. Guest column: inapproximability results via long code based pcps. *SIGACT News*, 36(2):25–42, 2005.
- [62] S. Khot. Inapproximability of NP-complete problems, discrete fourier analysis, and geometry. In *Proc. the International Congress of Mathematicians*, 2010.

- [63] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? In *Proc. 45th IEEE Symposium on Foundations of Computer Science*, pages 146–154, 2004.
- [64] S. Khot and A. Naor. Approximate kernel clustering. In *Proc. 49th IEEE Symposium on Foundations of Computer Science*, 2008.
- [65] S. Khot and A. Naor. Sharp kernel clustering algorithms and their associated Grothendieck inequalities. In *SODA*, 2010.
- [66] S. Khot and R. O’Donnell. Sdp gaps and ugc-hardness for maxcutgain. In *Proc. Annual IEEE Symposium on Foundations of Computer Science*, pages 217–226, 2006.
- [67] S. Khot, P. Popat, and R. Saket. Approximate Lasserre integrality gap for unique games. *Manuscript*, 2010.
- [68] S. Khot and O. Regev. Vertex cover might be hard to approximate to within $2 - \epsilon$. In *Proc. 18th IEEE Conference on Computational Complexity*, 2003.
- [69] S. Khot and R. Saket. A 3-query non-adaptive pcg with perfect completeness. In *IEEE Conference on Computational Complexity*, pages 159–169, 2006.
- [70] S. Khot and R. Saket. Sdp integrality gaps with local ℓ_1 -embeddability. In *Proc. 50th IEEE Symposium on Foundations of Computer Science*, pages 565–574, 2009.
- [71] S. Khot and N. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1 . In *Proc. 46th IEEE Symposium on Foundations of Computer Science*, 2005.
- [72] G. Kindler, A. Naor, and G. Schechtman. The ugc hardness threshold of the l_p Grothendieck problem. In *SODA*, pages 64–73, 2008.
- [73] G. Kindler, R. O’Donnell, A. Rao, and A. Wigderson. Spherical cubes and rounding in high dimensions. In *Proc. Annual IEEE Symposium on Foundations of Computer Science*, pages 189–198, 2008.
- [74] A. Kolla. Spectral algorithms for unique games. In *Proc. Annual IEEE Conference on Computational Complexity*, 2010.
- [75] R. Krauthgamer and Y. Rabani. Improved lower bounds for embeddings into l_1 . In *Proc. ACM-SIAM Symposium on Discrete Algorithms*, 2006.
- [76] J. R. Lee and A. Naor. l_p metrics on the Heisenberg group and the Goemans-Linial conjecture. In *Proc. 47th IEEE Symposium on Foundations of Computer Science*, pages 99–108, 2006.
- [77] F. T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, 1999.
- [78] M. Lewin, D. Livnat, and U. Zwick. Improved rounding techniques for the max 2-sat and max di-cut problems. In *IPCO*, pages 67–82, 2002.
- [79] N. Linial. Finite metric spaces-combinatorics, geometry and algorithms. In *Proc. International Congress of Mathematicians*, volume 3, pages 573–586, 2002.

- [80] K. Makarychev and Y. Makarychev. How to play unique games on expanders. *CoRR*, abs/0903.0367, 2009.
- [81] R. Manokaran, J. Naor, P. Raghavendra, and R. Schwartz. Sdp gaps and ugc hardness for multiway cut, 0-extension, and metric labeling. In *Proc. ACM Symposium on the Theory of Computing*, pages 11–20, 2008.
- [82] E. Mossel. Gaussian bounds for noise correlation of functions and tight analysis of long codes. In *Proc. Annual IEEE Symposium on Foundations of Computer Science*, pages 156–165, 2008.
- [83] E. Mossel, R. O’Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. In *Proc. 46th IEEE Symposium on Foundations of Computer Science*, pages 21–30, 2005.
- [84] A. Naor. On the banach space valued Azuma inequality and small set isoperimetry in Alon-Roichman graphs. *Unpublished manuscript*, 2004.
- [85] A. Naor. L1 embeddings of the Heisenberg group and fast estimation of graph isoperimetry. In *Proc. the International Congress of Mathematicians*, 2010.
- [86] A. Newman. *Approximating the maximum acyclic subgraph*. Masters’ Thesis, MIT, 2000.
- [87] R. O’Donnell and Y. Wu. An optimal sdp algorithm for max-cut, and equally optimal long code tests. In *Proc. ACM Symposium on the Theory of Computing*, pages 335–344, 2008.
- [88] R. O’Donnell and Y. Wu. Conditional hardness for satisfiable 3-csp. In *Proc. ACM Symposium on the Theory of Computing*, pages 493–502, 2009.
- [89] C. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and Systems Sciences*, 43:425–440, 1991.
- [90] P. Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proc. ACM Symposium on the Theory of Computing*, pages 245–254, 2008.
- [91] P. Raghavendra and D. Steurer. How to round any csp. In *Proc. Annual IEEE Symposium on Foundations of Computer Science*, pages 586–594, 2009.
- [92] P. Raghavendra and D. Steurer. Integrality gaps for strong sdp relaxations of unique games. In *Proc. 50th IEEE Symposium on Foundations of Computer Science*, pages 575–585, 2009.
- [93] P. Raghavendra and D. Steurer. Towards computing the Grothendieck constant. In *Proc. ACM-SIAM Symposium on Discrete Algorithms*, pages 525–534, 2009.
- [94] P. Raghavendra and D. Steurer. Graph expansion and the unique games conjecture. In *Proc. 42nd ACM Symposium on Theory of Computing*, 2010.
- [95] A. Rao. Parallel repetition in projection games and a concentration bound. In *Proc. ACM Symposium on the Theory of Computing*, pages 1–10, 2008.
- [96] R. Raz. A parallel repetition theorem. *SIAM J. of Computing*, 27(3):763–803, 1998.
- [97] R. Raz. A counterexample to strong parallel repetition. In *Proc. Annual IEEE Symposium on Foundations of Computer Science*, pages 369–373, 2008.

- [98] V. Rotar'. Limit theorems for polylinear forms. *J. Multivariate Anal.*, 9(4):511–530, 1979.
- [99] L. Russo. An approximate zero-one law. *Z. Wahrsch. Verw. Gebiete*, 61(1):129–139, 1982.
- [100] A. Samorodnitsky and L. Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *Proc. 32nd ACM Symposium on Theory of Computing*, pages 191–199, 2000.
- [101] A. Samorodnitsky and L. Trevisan. Gowers uniformity, influence of variables, and pcps. In *Proc. 38th ACM Symposium on Theory of Computing*, 2006.
- [102] P. Seymour. Packing directed circuits fractionally. *Combinatorica*, 15:281–288, 1995.
- [103] L. Trevisan. Approximation algorithms for unique games. In *Proc. 46th IEEE Symposium on Foundations of Computer Science*, 2005.
- [104] L. Trevisan. Inapproximability of combinatorial optimization problems. *Optimisation Combinatoire 2*, (Vangelis Paschos, Editor), Hermes, 2005.
- [105] U. Zwick. Approximation algorithms for constraint satisfaction problems involving at most three variables per constraint. In *SODA*, pages 201–210, 1998.