Optimal Inapproximability Results for MAX-CUT and Other 2-variable CSPs?

Subhash Khot[†]
School of Mathematics
Institute for Advanced Study
Princeton, NJ
khot@ias.edu

Guy Kindler
DIMACS
Rutgers University
Piscataway, NJ
gkindler@dimacs.rutgers.edu

Elchanan Mossel[*]
Department of Statistics
U.C. Berkeley
Berkeley, CA
mossel@stat.berkeley.edu

Ryan O'Donnell[†]
School of Mathematics
Institute for Advanced Study
Princeton, NJ
odonnell@ias.edu

Abstract

In this paper we give evidence suggesting that MAX-CUT is NP-hard to approximate to within a factor of $\alpha_{GW}+\epsilon$*, for all* $\epsilon > 0$*, where* α_{GW} *denotes the approximation ratio achieved by the Goemans-Williamson algorithm [14],* $\alpha_{GW} \approx .878567$*. This result is conditional, relying on two conjectures: a) the Unique Games conjecture of Khot [24]; and, b) a very believable conjecture we call the Majority Is Stablest conjecture. These results indicate that the geometric nature of the Goemans-Williamson algorithm might be intrinsic to the MAX-CUT problem.*

The same two conjectures also imply that it is NP-hard to $(\beta+\epsilon)$*-approximate MAX-2SAT, where* $\beta \approx .943943$ *is the minimum of* $(2+\frac{2}{\pi}\theta)/(3-\cos(\theta))$ *on* $(\frac{\pi}{2},\pi)$*. Motivated by our proof techniques, we show that if the MAX-2CSP and MAX-2SAT problems are slightly restricted — in a way that seems to retain all their hardness — then they have* $(\alpha_{GW}-\epsilon)$*- and* $(\beta-\epsilon)$*-approximation algorithms, respectively.*

Though we are unable to prove the Majority Is Stablest conjecture, we give some partial results and indicate possible directions of attack. Our partial results are enough to imply that MAX-CUT is hard to $(\frac{3}{4}+\frac{1}{2\pi}+\epsilon)$*-approximate (*$\approx .909155$*), assuming only the Unique Games conjecture. We also discuss MAX-2CSP problems over non-boolean domains and state some related results and conjectures. We show, for example, that the Unique Games conjecture implies that it is hard to approximate MAX-2LIN(q) to within any constant factor.*

1. Introduction

The main result in this paper is a bound on the approximability of the MAX-CUT problem which matches the approximation ratio achieved by the well-known Goemans-Williamson algorithm [14]. The proof of this hardness result unfortunately relies on two unproven conjectures. These conjectures are the *Unique Games conjecture* of Khot [24] and a commonly believed conjecture we call the *Majority Is Stablest conjecture*. For the convenience of the reader we will now briefly describe the conjectures; formal statements appear in Sections 3 and 4, respectively.

Unique Games conjecture (roughly): Given a bipartite graph G, a large constant size set of labels $[M]$, and a permutation of $[M]$ written on each edge, consider the problem of trying to find a labeling of the vertices of G from $[M]$ so that each edge permutation is 'satisfied.' The conjecture is that if M is a large enough constant then it is NP-hard to distinguish instances which are 99% satisfiable from instances which are 1% satisfiable.

Majority Is Stablest conjecture (roughly): Let f be a boolean function which is equally often 0 or 1. Suppose the string x is picked uniformly at random and the string y is formed by flipping each bit of x independently with probability η; we call $\Pr[f(x) = f(y)]$ the

∗ Supported by a Miller fellowship in Computer Science and Statistics, U.C. Berkeley

† This material is based upon work supported by the National Science Foundation under agreement Nos. DMS-0111298 and CCR-0324906 respectively. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

noise stability of f. The conjecture states that among all f in which each coordinate has $o(1)$ 'influence,' the Majority function has the highest noise stability, up to an additive $o(1)$.

We add in passing that the name Majority Is Stablest is a bit of a misnomer in that almost all balanced boolean threshold functions are equally noise stable (see Theorem 5).

Let us discuss the reasons why we believe this result is important despite its reliance on unproven conjectures. First, we believe it is quite remarkable that these two conjectures should yield a *matching* hardness of approximation ratio for MAX-CUT, and that indeed the best factor should be the peculiar number α_{GW}. It is intriguing that the precise quantity α_{GW} should arise from a noise stability property of the Majority function, and certainly there was previously little evidence to suggest that the Goemans-Williamson algorithm might be optimal.

As regards the conjectures themselves, we strongly believe in the Majority Is Stablest conjecture. The experts in "analysis of boolean functions" whom we have consulted have agreed that the conjecture should be correct (in fact, similar conjectures already appear in the literature, e.g. Conjecture 5.1 in [21]), and every relevant piece of evidence is in concordance with the conjecture. Because of this, we believe that understanding the status of the Unique Games conjecture is the main issue.

Unlike the Majority Is Stablest conjecture, the Unique Games conjecture is far from certain to be true; in fact, there is no particularly strong evidence either for it or against it. Rather than this being a problem for our studies of MAX-CUT, we can view this paper as an investigation into the Unique Games conjecture via the lens of MAX-CUT. First, we see that the conjecture does not give an incorrectly strong hardness bound for MAX-CUT, and indeed (as it does for Vertex Cover [25]) it gives what would ultimately be a natural bound. Second, we show (by reduction) that (modulo the Majority Is Stablest conjecture) the Unique Games problem is not harder than the problem of beating the Goemans-Williamson algorithm for MAX-CUT; thus we give encouragement for attacking Unique Games algorithmically.

Another reason we believe our result is important despite its reliance on conjectures is related to this last point. Since the Goemans-Williamson algorithm was published a decade ago there has been no algorithmic progress on approximating MAX-CUT. Since Håstad's classic inapproximability paper [17] from two years later there has been no progress on the hardness of approximating MAX-CUT, except for the creation of a better reduction gadget. As one of the most natural and simple problems to have resisted matching approximability bounds, we feel MAX-CUT deserves further investigation and analysis. In particular, we think that regardless of the truth of the Unique Games conjecture, this paper gives interesting insight into the geometric nature of MAX-CUT. Indeed, insights we have gleaned from studying the MAX-CUT problem in this light have motivated us to give new positive approximation algorithms for variants of other 2-variable CSPs such as MAX-2SAT.

Next, we believe that the conjectures and open questions we consider in this work, and in particular the Majority Is Stablest conjecture itself, are of significant independent interest. First, the conjecture has interesting applications outside of this work — to the economic theory of social choice [21] for example — and may well prove useful for other PCP-based inapproximability results. Second, the conjecture is an extension of, or is very similar to, several other important theorems in the analysis of boolean functions, including the KKL theorem [20] and Bourgain's theorem [3]; and furthermore, the partial progress we make on proving the Majority Is Stablest conjecture clarifies certain aspects of the papers of Friedgut, Kalai, and Naor [13] (cf. Theorem 8) and Talagrand [30] (cf. Theorem 10). We note that our partial progress lets us prove an inapproximability factor of $.909155$ for MAX-CUT assuming only the Unique Games conjecture; this is already stronger than the best known bound.

Finally, considering analogues of the Majority is Stablest conjecture for q-ary functions where $q > 2$, raises independently interesting questions, which are also relevant for hardness-of-approximation. This is discussed more thoroughly in the full version of the paper.

2. About MAX-CUT

The MAX-CUT problem is a classic and simple combinatorial optimization problem: Given a graph G, find the size of the largest cut in G. By a cut we mean a partition of the vertices of G into two sets; the size of the cut is the number of edges with one vertex on either side of the partition. One can also consider a weighted version of the problem in which each edge is assigned a nonnegative weight and the goal is to cut as much weight as possible.

MAX-CUT is NP-complete (indeed, it is one of Karp's original NP-complete problems [23]) and so it is of interest to try to find polynomial time approximation algorithms. For maximization problems such as MAX-CUT we say an algorithm gives an α-approximation if it always returns an answer which is at least α times the optimal value; we also often relax this definition to allow randomized algorithms which in expectation give α-approximations. Crescenzi, Silvestri, and Trevisan [6] have shown that the weighted and unweighted versions of MAX-CUT have equal optimal approximation factors (up to an additive $o(1)$) and so we pass freely between the two problems in this paper.

The trivial randomized algorithm for MAX-CUT — put each vertex on either side of the partition independently with equal probability — is a $1/2$-approximation, and this algorithm is easy to derandomize; Sahni and Gonzalez [27] gave the first $1/2$-approximation algorithm in 1976. Following this some $(1/2 + o(1))$-approximation algorithms were given, but no real progress was made until the breakthrough 1994 paper of Goemans and Williamson [14]. This remarkable work used semidefinite programming to achieve an α_{GW}-approximation algorithm, where the constant $\alpha_{\text{GW}} \approx .878567$ is the trigonometric quantity

$$\alpha_{\text{GW}} = \min_{0<\theta<\pi} \frac{\theta/\pi}{(1-\cos\theta)/2}.$$

The optimal choice of θ is the solution of $\theta = \tan(\theta/2)$, namely $\theta^* \approx 2.33 \approx 134^\circ$, and $\alpha_{\text{GW}} = \frac{2}{\pi \sin\theta^*}$. The geometric nature of Goemans and Williamson's algorithm might be considered surprising, but as we shall see, this geometry seems to be an inherent part of the MAX-CUT problem.

On the hardness of approximation side, MAX-CUT was proved MAX-SNP hard [26] and Bellare, Goldreich, and Sudan [1] explicitly showed that it was NP-hard to approximate MAX-CUT to any factor higher than $83/84$. The hardness factor was improved to $16/17 \approx .941176$ by Håstad [19] via a reduction from MAX-3LIN using a gadget of Trevisan, Sorkin, Sudan, and Williamson [31]. This stands as the current best hardness result.

Despite much effort and many improvements in the approximation guarantees of other semidefinite programming-based algorithms, no one has been able to improve on the algorithm of Goemans and Williamson. Although the true approximation ratio of Goemans-Williamson was proved to be not more than α_{GW} [22, 11] and the integrality gap of their semidefinite relaxation was also proved to be α_{GW} [11], there appears on the face of it to be plenty of possibilities for improvement. Adding triangle constraints and other valid constraints to the semidefinite program has been suggested, alternate rounding schemes have been proposed, and local modification heuristics that work for special graphs have been proven (see, e.g., [14, 9, 8, 22, 32, 10, 11]). And of course, perhaps a completely different algorithm altogether can perform better. Several papers have either explicitly ([8]) or implicitly ([11]) given the problem of improving on α_{GW} as an important research goal.

In this paper we give evidence that indeed it may be that MAX-CUT is hard to approximate within any factor larger than α_{GW}.

3. About the Unique Games conjecture

MAX-CUT belongs to the class of constraint satisfaction problems on 2 variables (2-CSPs). In a k-CSP we are given a set of variables and a set of constraints, where each constraint depends on exactly k variables. The goal is to find an assignment to the variables so as to maximize the number of constraints satisfied. In case of MAX-CUT, the vertices serve as variables and the edges as constraints. Every constraint says that two certain variables should receive different boolean values.

Proving inapproximability results for a k-CSP is equivalent to constructing a k-query PCP with a specific acceptance predicate. Usually the so-called Label Cover problem is a starting point for any PCP construction. Label Cover is a 2-CSP where the variables range over a large (non-boolean) domain. Usually, inapproximability results for boolean CSPs are obtained by encoding assignments to Label Cover variables via a binary code and then running PCP tests on the (supposed) encodings. This approach has been immensely successful in proving inapproximability results for k-CSPs with $k \geq 3$ (see for example [19, 28, 16]). However the approach gets stuck in the case of 2-CSPs. We seem to have no techniques for constructing boolean 2-query PCPs and the bottleneck seems to be the lack of an appropriate PCP 'outer verifier'.

Khot suggested the Unique Games Conjecture in [24] as a possible direction for proving inapproximability results for some important 2-CSPs, such as Min-2SAT-Deletion, Vertex Cover, Graph-Min-Bisection and MAX-CUT. This conjecture asserts the hardness of the 'Unique Label Cover' problem:

Definition 1. *The Unique Label Cover problem, $\mathcal{L}(V, W, E, [M], \{\pi^{v,w}\}_{(v,w)\in E})$ is defined as follows: Given is a regular bipartite graph with left side vertices V, right side vertices W, and a set of edges E. The goal is to assign one 'label' to every vertex of the graph, where $[M]$ is the set of allowed labels. The labeling is supposed to satisfy certain constraints given by bijective maps $\sigma_{v,w} : [M] \to [M]$. There is one such map for every edge $(v,w) \in E$. A labeling 'satisfies' an edge (v,w) if*

$$\sigma_{v,w}(\text{label}(w)) = \text{label}(v).$$

The optimum OPT of the unique label cover problem is defined to be the maximum fraction of edges satisfied by any labeling.

The Unique Label Cover problem is a special case of the Label Cover problem. It can also be stated in terms of 2-Prover-1-Round Games, but the Label Cover formulation is easier to work with. The Unique Games conjecture asserts that this problem is hard:

Unique Games conjecture: *For any η, $\delta > 0$, there exists a constant $M = M(\eta, \delta)$ such that it is NP-hard to distinguish whether the Unique Label Cover problem with label set of size M has optimum at least $1 - \eta$ or* *at most δ.*

The Unique Games conjectures asserts the existence of a powerful outer verifier that makes only 2 queries (albeit over a large alphabet) and has a very specific acceptance predicate: for every answer to the first query, there is exactly one answer to the second query for which the verifier would accept, and vice versa. Once we have such a powerful outer verifier, we can possibly construct a suitable inner verifier and prove the desired inapproximability results. However the inner verifier typically relies on rather deep theorems about the Fourier spectrum of boolean functions, e.g. the theorem of Bourgain [3] or of Friedgut [12].

The Unique Games conjecture was used in [24] to show that Min-2SAT-Deletion is NP-hard to approximate within any constant factor. The inner verifier is based on a test proposed by Håstad [18] and on Bourgain's theorem. Khot and Regev [25] showed that the conjecture implies that Vertex Cover is NP-hard to approximate within any factor less than 2. The inner verifier in their paper is based on Friedgut's theorem and is inspired by the work of Dinur and Safra [7] that showed 1.36 hardness for Vertex Cover. In the present paper we continue this line of research and propose a plausible direction for attacking the MAX-CUT problem. We do construct an inner verifier, but to prove its correctness we need another powerful conjecture about the Fourier spectrum of boolean functions. This is the subject of the Majority Is Stablest conjecture.

4. About the Majority Is Stablest conjecture

To state the Majority Is Stablest conjecture, we need some definitions. For convenience we regard the boolean values as -1 and 1 rather than 0 and 1. Thus a boolean function is a map $f\colon \{-1,1\}^n \to \{-1,1\}$. We will often generalize to the case of functions $f\colon \{-1,1\}^n \to \mathbb{R}$. In all of what follows we consider the set of strings $\{-1,1\}^n$ to be a probability space under the uniform distribution.

First we recall the well-known notion of 'influence', introduced to computer science in [2] and studied even earlier in economics.

Definition 2. *Let $f\colon \{-1,1\}^n \to \mathbb{R}$. Then the* influence of x_i on f *is defined by*

$$\mathrm{Inf}_i(f) = \mathop{\mathbf{E}}_{(x_1,\ldots,x_{i-1},x_{i+1},\ldots,x_n)} \left[\mathrm{Var}_{x_i}[f]\right].$$

(Note that for $f\colon \{-1,1\}^n \to \{-1,1\}$,

$$\mathrm{Inf}_i(f) = \Pr_{x\in\{-1,1\}^n}[f(x) \neq f(x_1,\ldots,-x_i,\ldots x_n)].)$$

Instead of picking x at random, flipping one bit, and seeing if this changes the value of f, we can instead flip a constant fraction (in expectation) of the bits. This leads to the study of 'noise sensitivity', pioneered in computer science by [20, 17, 4].

Definition 3. *Let $f\colon \{-1,1\}^n \to \mathbb{R}$ and let $-1 \le \rho \le 1$. The* noise correlation of f at ρ *is defined as follows: Let x be a uniformly random string in $\{-1,1\}^n$ and let y be a 'ρ-correlated' copy; i.e., pick each bit y_i independently so that $\mathbf{E}[x_i y_i] = \rho$. Then the noise correlation is defined to be*

$$\mathbb{S}_\rho(f) = \mathbf{E}_{x,y}[f(x)f(y)].$$

(Note that for $f\colon \{-1,1\}^n \to \{-1,1\}$, $\mathbb{S}_\rho(f) = 2\Pr_{x,y}[f(x) = f(y)] - 1$.)

We may now state the Majority Is Stablest conjecture. Informally, the conjecture says that among all balanced boolean functions with small influences, the Majority function has the highest noise correlation. Note that the assumption of small influences is necessary since the 'dictator' function $f(x) = x_i$ provably has the highest noise correlation among all balanced boolean functions, for every ρ. Note that when n tends to infinity, the noise correlation at ρ of the n-bit Majority function approaches $(1 - \frac{2}{\pi}\arccos\rho)$ (this fact was stated in a paper of Gulibaud from the 1960's [15] and is ultimately derived from the Central Limit theorem plus a result from an 1890's paper of Sheppard [29]). Thus we have the formal statement of the conjecture:

Majority Is Stablest conjecture: *Fix $\rho \in [0,1)$. Then for any $\epsilon > 0$ there is a small enough $\delta = \delta(\epsilon,\rho) > 0$ such that if $f : \{-1,1\}^n \to [-1,1]$ is any function satisfying $\mathbf{E}[f] = 0$ and $\mathrm{Inf}_i(f) \le \delta$ for all $i = 1\ldots n$, then*

$$\mathbb{S}_\rho(f) \le 1 - \tfrac{2}{\pi}\arccos\rho + \epsilon.$$

Regarding the plausibility of the conjecture: We strongly believe that the Majority Is Stablest conjecture is true, as do the experts in the field whom we consulted. We note that similar, though weaker conjectures already appear in the literature (see e.g. Section 5 in [21]). Further discussion of the plausibility of the conjecture appears in the full version of the paper.

In the remainder of this section, we shall describe why the Majority Is Stablest conjecture is relevant for MAX-CUT inner verifiers.

As described in the previous section, inapproximability results for many problems are obtained by constructing a tailor-made PCP; usually, the PCP is obtained by composing an 'outer verifier' (almost always a Label Cover problem) with an 'inner verifier'. As mentioned the outer verifier for our reduction is the Unique Label Cover problem. As for the inner verifier, it is always application-specific and its acceptance predicate is tailor-made for the problem at hand, in our case MAX-CUT.

A *codeword test* is an essential submodule of an inner verifier. It is a probabilistic procedure for checking whether a given string is a codeword of an error-correcting code, most commonly the 'Long Code' (see [1]).

Definition 4. *The Long Code over domain $[n]$ is a binary code in which the message space is in fact the set of truth tables of boolean functions $f : \{-1,1\}^n \to \{-1,1\}$. The codeword encoding the 'message' $i \in [n]$ is given by the ith dictator function; i.e., the function $f(x_1, x_2, \dots, x_n) = x_i$.*

A codeword test for the Long Code can often be extended to a full-fledged inner verifier. So in the following, we will focus only on a Long Code test. The choice of the test is determined by the problem at hand, in our case MAX-CUT. The test must read two bits from a Long Code and accept if and only if the values read are distinct. Note that a legal Long Code word, i.e. a dictator, is the truth table of a boolean function in which one coordinate has influence 1. Let us say that a function f is *far from being a Long Code* if all the coordinates have $o(1)$ influences (note that this is not a standard notion of being far from a codeword, but rather a notion tailored for our proof technique).

We expect the following from a codeword test: a correct Long Code word passes the test with probability c (called the 'completeness' parameter of the test) whereas any function far from being a Long Code passes the test with probability at most s (called the 'soundness' parameter). Once we construct a full-fledged inner verifier, the ratio s/c will be the inapproximability factor for MAX-CUT.

The Long Code test. As mentioned before, our Long Code test will need to take a boolean function $f : \{-1,1\}^n \to \{-1,1\}$, pick two inputs x and y, and check that $f(x) \neq f(y)$. In fact our test will be precisely a 'noise correlation' test for some fixed noise rate ρ; i.e., x will be chosen uniformly at random and y will be formed by flipping each bit of x independently with probability $(1-\rho)/2$. Here ρ will be a value between -1 and 0, and therefore y is a *highly* noisy version of x, or alternatively, a moderately noisy version of $-x$. Thus (at least for legal Long Code words) we expect $f(x)$ to be quite *anticorrelated* with $f(y)$; i.e., it should pass the test with relatively high probability. Recalling Definition 3, we see that the probability a given function f passes our test is precisely $\frac{1}{2} - \frac{1}{2}\mathbb{S}_\rho(f)$.

A legal Long Code word, i.e. a dictator function, has noise correlation precisely ρ and thus the completeness of the Long Code test is $c = \frac{1}{2} - \frac{1}{2}\rho$. The crucial aspect of our test is the analysis of the soundness parameter.

This is where the Majority Is Stablest conjecture comes in. Suppose $f\colon \{-1,1\}^n \to \{-1,1\}$ is any function that is far from being a Long Code word. By a simple trick (shown in the full version of the paper) we can show that the Majority Is Stablest conjecture (which is stated only for $\rho \geq 0$) implies that for $\rho < 0$ the noise correlation of f at ρ is at *least* $1 - \frac{2}{\pi}\arccos\rho$ (a negative number). Hence it follows that functions that are far from being a Long Code pass the test with probability *at most* $s = \frac{1}{2} - \frac{1}{2}(1 - \frac{2}{\pi}\arccos\rho) = (\arccos\rho)/\pi$.

Choosing $\rho < 0$ as we please, this leads to an inapproximability ratio of

$$\frac{s}{c} = \min_{-1<\rho<0} \frac{(\arccos\rho)/\pi}{\frac{1}{2} - \frac{1}{2}\rho} = \min_{0\leq\theta\leq\pi} \frac{\theta/\pi}{(1-\cos\theta)/2} = \alpha_{\text{GW}},$$

precisely the Goemans-Williamson constant.

5. On the geometry of MAX-CUT

We shall now try to explain (non-rigorously) the connection between the Majority Is Stablest conjecture and the geometric picture that arises from the Goemans-Williamson algorithm. But before going further, let us first note that the approximation ratio achieved by Goemans-Williamson arises as the solution of a trigonometric minimization problem, which in turn originates from a geometric setting. To obtain a matching inapproximability constant, it seems essential to introduce some similar geometric structure. Such a structure is present in the construction of our Long Code test, although it is only implicit in the actual proofs.

For the purposes of the following explanation, let us consider the n-dimensional discrete cube $\{-1,1\}^n$ as a subset of the n-dimensional Euclidean unit sphere (we normalize the Euclidean norm accordingly). The Majority Is Stablest conjecture essentially states that the discrete cube is a good approximation of the sphere in a certain sense.

The Goemans-Williamson algorithm. We start with a brief description of how the approximation ratio α_{GW} arises in the Goemans-Williamson algorithm. To find a large cut in a given graph $G = (V,E)$ with n vertices, the Goemans-Williamson algorithm embeds the graph in the unit sphere of $\mathbb{R}^n$, identifying each vertex $v \in V$ with a unit vector $\mathbf{x}_v$ on the sphere. The embedding is selected such that the sum

$$\sum_{(u,v)\in E} \frac{1}{2} - \frac{1}{2}\langle \mathbf{x}_u, \mathbf{x}_v \rangle, \tag{1}$$

involving the inner products of vectors associated with the endpoints of edges of G, is maximized. The maximal sum bounds from above the size of the maximum cut, since the size of every cut can be realized by associating all the vertices from one side of the cut with an arbitrary point $\mathbf{x}$ on the sphere, and associating all other vertices with $-\mathbf{x}$.

Once the embedding is set, a cut in G is obtained by choosing a random hyperplane through the origin and partitioning the vertices according to the side of the hyperplane on which their associated vectors fall. For an edge (u, v) in G, the probability that u and v lie on opposite sides of the random cut is proportional to the angle between $\mathbf{x}_u$ and $\mathbf{x}_v$. More precisely, letting $\rho = \langle \mathbf{x}_u, \mathbf{x}_v \rangle$ denote the inner product between the vectors associated with u and v, the probability that the edge (u, v) is cut is $(\arccos \rho)/\pi$.

The approximation ratio α_{GW} of the Goemans-Williamson algorithm is obtained by noting that

$$\alpha_{\text{GW}} = \min_{-1 \le \rho \le 1} \frac{(\arccos \rho)/\pi}{\frac{1}{2} - \frac{1}{2}\rho} \approx .878567 \qquad (2)$$

is the smallest ratio possible between the probability of an edge being cut and its contribution to (1). Hence the expected size of the cut obtained by the Goemans-Williamson algorithm is at least an α_{GW}-fraction of (1), and therefore it is also at least an α_{GW}-fraction of the maximum cut in G.

Cutting the sphere. In [11], Feige and Schechtman considered the graph G_ρ whose vertices are all the vectors on the unit sphere and in which two vertices are connected by an edge in G_ρ iff their inner product is roughly ρ (we do not get into the precise details). It is shown in [11] that in this graph the largest cut is obtained by any hyperplane through the origin. (To state this rigorously one should define appropriate measures etc., but let us remain at a simplistic level for now.) Such a hyperplane cuts an $(\arccos \rho)/\pi$-fraction of the edges in the graph.

Restricting to the cube. We would like to consider an edge-weighted graph H_ρ which is, in a non-rigorous sense, the graph induced by G_ρ on the discrete hypercube. For two vectors $\mathbf{x}, \mathbf{y}$ on the discrete cube, we define the weight of the edge $(\mathbf{x}, \mathbf{y})$ to be

$$\Pr[X = \mathbf{x} \text{ and } Y = \mathbf{y}],$$

where X and Y are ρ-correlated random elements of the discrete cube. The graph H_ρ resembles G_ρ in the sense that almost all the edge-weight in H_ρ is concentrated on edges $(\mathbf{x}, \mathbf{y})$ for which $\langle \mathbf{x}, \mathbf{y} \rangle \approx \rho$; we call such edges *typical edges*. Let us examine how good H_ρ is as an 'approximation' of the graph G_ρ.

Note that the structure of H_ρ is very reminiscent of our Long Code test, mentioned above. To make the similarity even clearer, note that a cut C in H_ρ immediately defines a boolean function f_C over the discrete cube. It is easy to observe that the size of C (namely the sum of weights of the edges that are cut) is exactly the noise stability of f — i.e., the acceptance probability of the Long Code test with parameter ρ when applied to f_C.

The size of the cut. So how large can the size of C be? If C is determined by a random hyperplane, then a typical edge is cut with probability about $(\arccos \rho)/\pi$. The expected size of such a cut is therefore roughly the same as the weight of the maximal cut in G_ρ (when the total weight of the edges in G_ρ is normalized to 1).

There are, however, cuts in H_ρ whose weight is larger than $(\arccos \rho)/\pi$. For example, one can partition the vertices in H_ρ according to their first coordinate, taking one side of the cut C to be the set of vectors in the discrete cube whose first coordinate is 1 and the other side of C to be the set of vectors whose first coordinate is -1; note that this is the cut defined by the hyperplane which is perpendicular to the first coordinate. When interpreted as a function, C corresponds to the function $f_C(x) = x_1$; i.e., it is a correct Long Code word. One can easily observe that the size of C is $\frac{1}{2} - \frac{1}{2}\rho$ — i.e., it is exactly the completeness of the Long Code test with parameter ρ.

The conjecture comes in. The size of one-coordinate cuts in H_ρ is larger than the best cuts achievable in G_ρ. The Majority Is Stablest conjecture implies, however, that essentially those are the only special cases, and that all other cuts in H_ρ are no larger than the maximum cut in G_ρ. That is, it implies that unless f_C depends significantly on one of the coordinates, then the size of C is at most $(\arccos \rho)/\pi + \epsilon$. This is stated formally in the following proposition:

Proposition *If the Majority Is Stablest conjecture is true, then the following holds for every* $\rho \in (-1, 0]$. *For any* $\epsilon > 0$ *there is a small enough* $\delta = \delta(\epsilon, \rho) > 0$ *such that if* C *is a cut in* H_ρ *such that* $\mathrm{Inf}_i(f_C) \le \delta$ *for every* i*, then the size of* C *is at most* $(\arccos \rho)/\pi + \epsilon$

In the full version of the paper we prove that the statement of the above Proposition holds with respect to all *hyperplane cuts*, even without assuming the Majority Is Stablest conjecture.

6. Our results

We now formally state our main results. All of these will be proven in the full version of this paper.

6.1. Hardness for Max-Cut and 2-bit CSPs

Our main results regarding MAX-CUT are the following:

Theorem 1. *Assume the Unique Games conjecture and the Majority Is Stablest conjecture. Then it is NP-hard to approximate MAX-CUT to within any factor greater than the Goemans-Williamson constant,* α_{GW}*, where*

$$\alpha_{GW} = \min_{0 \le \theta \le \pi} \frac{\theta/\pi}{(1 - \cos\theta)/2} = \min_{-1 < \rho < 0} \frac{(\arccos \rho)/\pi}{(1 - \rho)/2}.$$

Theorem 2. *Assume only the Unique Games conjecture. Then it is NP-hard to approximate MAX-CUT to within any factor greater than* $3/4 + 1/2\pi$.

In the full version of the paper we also discuss implications of our results for other 2-bit CSPs besides MAX-CUT. In particular we prove:

Theorem 3. *Assume the Unique Games conjecture and the Majority Is Stablest conjecture. Then it is NP-hard to approximate MAX-2SAT to within any factor greater than* $\beta \approx .943943$*, where*

$$\beta = \min_{\frac{\pi}{2} \le \theta \le \pi} \frac{2 + (2/\pi)\theta}{3 - \cos\theta}.$$

The proof of Theorem 3 actually implies that MAX-2SAT is hard to approximate to within any factor greater than β, even if restricted to instances where each variable appears equally often positively and negatively. We show that for this restricted problem, called Balanced-MAX-2SAT, the approximation bound β is tight; i.e., it *can* be approximated to within any factor smaller than β:

Theorem 4. *Balanced-MAX-2SAT is polynomial-time approximable to within any factor smaller than* β.

6.2. Partial progress on the Majority Is Stablest conjecture

We now state some partial results in the direction of the Majority Is Stablest conjecture.

First, as mentioned we have shown that the conjecture holds for the subclass of balanced threshold functions; this follows from the following two results:

Theorem 5. *Let* $f : \{-1,1\}^n \to \{-1,1\}$ *be any balanced threshold function,* $f(x) = \mathrm{sgn}(a_1x_1 + \cdots a_nx_n)$[1]*, where* $\sum a_i^2 = 1$*. Let* $\delta = \max\{|a_i|\}$*. Then for all* $\rho \in [-1,1]$*,*

$$\mathbb{S}_\rho(f) = 1 - \tfrac{2}{\pi}\arccos\rho \pm O(\delta(1-|\rho|)^{-3/2}).$$

Proposition 6. *Let* $f : \{-1,1\}^n \to \{-1,1\}$ *be any balanced threshold function,* $f(x) = \mathrm{sgn}(a_1x_1 + \cdots a_nx_n)$*, where* $\sum a_i^2 = 1$*. Then for* $\delta = \max\{|a_i|\}$*, it holds that* $\max_i \{\mathrm{Inf}_i(f)\} \ge \Omega(\delta)$*.*

Some of our partial results are stated in terms of Fourier representation. To clarify the relevance of the Fourier representation to noise sensitivity of boolean functions, let us recall the well-known expression for noise correlation in terms of Fourier representation:

Proposition 7. *Let* $f\colon \{-1,1\}^n \to \mathbb{R}$*. Then for every* $\rho \in [-1,1]$*,* $\mathbb{S}_\rho(f) = \sum_{S\subseteq[n]} \rho^{|S|}\hat{f}(S)^2$ *.*

1 Without loss of generality we assume the linear form is never 0.

Since for a function $f : \{-1,1\}^n \to [-1,1]$, the sum $\sum_S \hat{f}(S)^2$ is bounded by 1, the stability of such a function is maximized by putting as much weight as possible on the low-level coefficients (by the 'level' of a coefficient $\hat{f}(S)$ we mean the cardinality of S). The weight of a balanced function f on the zero level must, however, be zero. We can also prove that any function with small influences has no more Fourier weight at level 1 than the majority function.

Theorem 8. *Suppose* $f : \{-1,1\}^n \to [-1,1]$ *satisfies* $\mathrm{Inf}_i(f) \le \delta$ *for all* i*. Then* $\sum_{|S|=1} \hat{f}(S)^2 \le \frac{2}{\pi} + C\delta$ *, where* $C = 2(1-\sqrt{2/\pi})$*.*

Using Theorem 8 together with Proposition 7, we get the following simple corollary, which we view as a weakened version of the Majority Is Stablest conjecture.

Corollary 9. *Suppose* $f\colon \{-1,1\}^n \to [-1,1]$ *satisfies* $\mathrm{Inf}_i(f) \le \delta$ *for all* i*, and assume* f *is balanced, namely* $\mathbf{E}[f] = 0$*. Then for* $C = 2(1-\sqrt{2/\pi})$ *it holds that*

$$\mathbb{S}_\rho(f) \le \tfrac{2}{\pi}\rho + (1-\tfrac{2}{\pi})\rho^2 + C\delta(\rho - \rho^2).$$

Theorem 2 is obtained by plugging Corollary 9 instead of the Majority Is Stablest conjecture, in the proof of Theorem 1. To get a tighter form of Corollary 9 (and therefore stronger parameters in Theorem 2), it may be helpful to get an improved bound in Theorem 8 for unbalanced functions: known techniques in bounding Fourier weights at low levels first apply random restrictions, and then bound the Fourier weight at level 1 of the resulting function; c.f. [5, 30, 3]. When performing a random restriction of a balanced function, the resulting function may be unbalanced.

We give the following generalization of Theorem 8, which essentially also generalizes the theorem of Talagrand [30] which states that for every function $f : \{-1,1\}^n \to \{-1,1\}$ with $\Pr[f = 1] = p \le 1/2$ it holds that $\sum_{|S|=1} \hat{f}(S)^2 \le O(p^2\log(1/p))$.

Theorem 10. *Let* ϕ *be the Gaussian density function and* Φ *be the Gaussian distribution function. Let* $U(x) = \phi(\Phi^{-1}(x)) : [0,1] \to [0, 1/\sqrt{2\pi}]$ *denote the so-called 'Gaussian isoperimetric function'.*

Suppose $f : \{-1,1\}^n \to [-1,1]$ *satisfies* $\mathrm{Inf}_i(f) \le \delta$ *for all* i*. Letting* $\mu = \frac{1}{2} + \frac{1}{2}\mathbf{E}[f]$*, we have*

$$\sum_{|S|=1} \hat{f}(S)^2 \le 4\left(U(\mu) + \epsilon\right)^2,$$

where the error term ϵ *is given by*

$$\epsilon = \max\{1, \sqrt{|\Phi^{-1}(\mu)|}\}O(\sqrt{\delta}).$$

6.3. Larger domains: q-ary functions

For a positive integer q, the problem MAX-2LIN(q) is that of maximizing the number of satisfied equations in a given system of linear equations modulo q, where exactly two variables appear in each equation. Note that MAX-CUT instances can also be regarded as instances of MAX-2LIN(2), where nodes correspond to variables, and each edge $e = (u, v)$ corresponds to the equation $u + v = 1 \pmod 2$.

While the Majority Is Stablest conjecture has applications to the MAX-CUT problem, results and conjectures concerning the noise stability of q-*ary* functions of the form $f : [q]^n \to [q]$ have applications to the hardness of MAX-2LIN(q).

Definition 5. *Let $f : [q]^n \to [q]$, and let x be a uniformly distributed element in $[q]^n$. The* noise stability of f at ρ *is defined by $\Pr_{x,y}[f(x) = f(y)]$, where each coordinate y_i of y is independently chosen to be equal to x_i with probability $(1-\rho)$, or to be uniformly distributed in $[q]$ with probability ρ.*

For $1 \le i \le n$, the influence *of the ith coordinate on f is defined by $\mathrm{Inf}_i(f) = \Pr_{x,x'}[f(x) \neq f(x')]$, where x' is obtained by replacing the ith coordinate of x with a uniformly chosen element in $[q]$.*

A q-ary function f is called *balanced* if it obtains each element $a \in [q]$ equally often. We conjecture that the stability of balanced q-ary functions with $o(1)$ influences must tend to zero as q grows.

Conjecture 11. *Let $0 < \rho < 1$ be any fixed parameter. Then there exist positive functions $\delta_\rho, S_\rho : \mathbb{N} \to \mathbb{R}$ such that $\lim_{q\to\infty} S_\rho(q) = 0$ and such that for every balanced relaxed q-ary function $f : [q]^n \to [q]$, all of whose influences are bounded by $\delta_\rho(q)$, the noise stability of f at ρ is at most $S_\rho(q)$.*

Conjecture 11 leads to an inapproximability result for MAX-2LIN(q).

Theorem 12. *The Unique Games conjecture and Conjecture 11 together imply the following. Let $\epsilon > 0$ be any fixed parameter. Then there exists a large enough q, such that given an instance of MAX-2LIN(q) is it NP-hard to distinguish between the case where it is ϵ-satisfiable and the case where it is $(1-\epsilon)$-satisfiable.*

We do not have an exact conjecture stating a precise bound on the noise stability of balanced q-ary functions. There is, however, a generalization of majority to the q-ary domain, which seems like a natural candidate for being the the stablest q-ary function. This is the *plurality* function, whose output on $x \in [q]^n$ is the most commonly appearing coordinate value in x. When $q \to \infty$ with n 'unbounded', we have an asymptotically sharp formula for the noise stability of plurality.

Theorem 13. *The noise stability of the plurality function at ρ for $n = \infty$ and $q \to \infty$ is $q^{-(1-\rho)/(1+\rho)+o(1)}$.* (Note that for any constant $\rho < 1$, this approaches 0 as q tends to ∞.)

Real-valued functions. Every real-valued function over $[q]^n$, $f : [q]^n \to \mathbb{R}$, can be written in a unique way in a "generalized Fourier expansion", as a sum $f = \sum_{S\subset[n]} f_S$, where each function $f_S(x)$ depends only on the coordinates of S and where $\mathbf{E}[f_S f_T] = 0$ for every $S \neq T$. In the spirit of Proposition 7, we can define the noise stability of such a function by $\sum_S \rho^{|S|}\|f_S\|_2^2$, and its ith influence by $\sum_{S:i\in S}\|f_S\|_2^2$. Under these definitions, the following conjecture about real-valued functions implies Conjecture 11.

Conjecture 14. *Let ρ, $-1 < \rho < 1$ be some fixed parameter. Then there exist positive functions $\delta_\rho, C_\rho : \mathbb{N} \to \mathbb{R}$ such that $\lim_{q\to\infty} C_\rho[q] = 0$ and such that the following holds. For every function $f : [q]^n \to [0,1]$ with $\mathbf{E}[f] = 1/q$, all of whose influences are smaller than $\delta_\rho(q)$,*

$$\sum_{S\neq\emptyset} \rho^{|S|}\|f_S\|_2^2 \le C_\rho(q)/q.$$

We have some partial results in the direction of Conjecture 14, bounding the first-level weight in the generalized Fourier representation of balanced q-ary functions. This bound, together with the Unique Games conjecture, already implies that MAX-2LIN(q) is NP-hard to approximate within any constant factor.

Theorem 15. *There exists a constant K and a positive function $\delta : \mathbb{N} \to \mathbb{R}$, such that for all functions $f : [q]^n \to [0,1]$ with $\mathbf{E}[f] \le 1/q$ and which satisfy $\|f_{\{i\}}\|_2^2 \le \delta(q)$ for all i, it holds that*

$$\sum_i \|f_{\{i\}}\|_2^2 \le \frac{K \log q}{q^2}.$$

Theorem 16. *Assuming the Unique Games conjecture, the following holds: For every $\epsilon > 0$ there exists a positive δ and an integer q such that given an instance of MAX-2LIN(q), it is NP-hard to distinguish between the case where the instance is δ-satisfiable and the case where it is $\epsilon\delta$-satisfiable.*

7. Acknowledgements

We thank Jean Bourgain, Steve Evans, Elad Hazan, Gil Kalai, Wenbo Li and Gideon Schechtman for stimulating discussion.

References

[1] M. Bellare, O. Goldreich, and M. Sudan. Free bits, PCPs and non-approximability — towards tight results. *SIAM Journal on Computing*, 27:804–915, 1998.

[2] M. Ben-Or and N. Linial. Collective coin flipping. In S. Micali, editor, *Randomness and Computation*. Academic Press, New York, 1990.

[3] J. Bourgain. On the distribution of the Fourier spectrum of boolean functions. *Israel Journal of Mathematics*, 131:269–276, 2002.

[4] I. Benjamini, G. Kalai, and O. Schramm. Noise sensitivity of boolean functions and applications to percolation. *Inst. Hautes Études Sci. Publ. Math.*, 90:5–43, 1999.

[5] J. Bourgain and G. Kalai. Influences of variables and threshold intervals under group symmetries. *Geom. and Func. Analysis*, 7:438–461, 1997.

[6] P. Crescenzi, R. Silvestri, and L. Trevisan. On weighted vs unweighted versions of combinatorial optimization problems. *Information and Computation*, 167(1):10–26, 2001.

[7] I. Dinur and S. Safra. The importance of being biased. In *Proc. 34th Ann. ACM Symp. on the Theory of Computing*, pages 33–42, 2002.

[8] U. Feige. *Randomized rounding of semidefinite programs — variations on the MAX-CUT example*, volume 1761 of *Lecture Notes in Computer Science*, pages 189–196. 1999.

[9] U. Feige and M. Goemans. Approximating the value of two prover proof systems, with applications to MAX-2SAT and MAX-DICUT. In *Proc. 3rd Ann. Israel Symp. on Theory of Comp. and Sys.*, pages 182–189, 1995.

[10] U. Feige, M. Karpinski, and M. Langberg. Improved approximation of MAX-CUT on graphs of bounded degree. *Electronic Colloq. on Comp. Complexity (ECCC)*, (021), 2000.

[11] U. Feige and G. Schechtman. On the optimality of the random hyperplane rounding technique for MAX-CUT. *Random Structures and Algorithms*, 20:403–440, 2002.

[12] E. Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):474–483, 1998.

[13] E. Friedgut, G. Kalai, and A. Naor. Boolean functions whose Fourier transform is concentrated at the first two levels. To appear, *Advances in Mathematics*.

[14] M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42:1115–1145, 1995.

[15] G. Guilbaud. *Theories of general interest, and the logical problem of aggregation*, pages 262–307. MIT PRess, 1966.

[16] V. Guruswami, J. Håstad, and M. Sudan. Hardness of approximate hypergraph coloring. In *Proc. 41st Ann. IEEE Symp. on Foundations of Comp. Sci.*, pages 149–158, 2000. To appear, SIAM Journal on Computing.

[17] J. Håstad. Some optimal inapproximability results. In *Proc. 29th Ann. ACM Symp. on the Theory of Computing*, pages 1–10, 1997.

[18] J. Håstad. On a protocol possibly useful for MIN-2SAT. Unpublished manuscript, 2001.

[19] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48:798–869, 2001.

[20] J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *Proc. 29th Ann. IEEE Symp. on Foundations of Comp. Sci.*, pages 68–80, 1988.

[21] G. Kalai. A Fourier-theoretic perspective on the Concordet paradox and Arrow's theorem. *Adv. in Appl. Math.*, 29(3):412–426, 2002.

[22] H. Karloff. How good is the Goemans-Williamson MAX-CUT algorithm? *SIAM Journal on Computing*, 29(1):336–350, 1999.

[23] R. Karp. *Reducibility among combinatorial problems*, pages 85–103. Plenum Press, 1972.

[24] S. Khot. On the power of unique 2-prover 1-round games. In *Proc. 34th Ann. ACM Symp. on the Theory of Computing*, pages 767–775, 2002.

[25] S. Khot and O. Regev. Vertex Cover might be hard to approximate to within $2-\epsilon$. In *Proc. 18th Ann. IEEE Confence on Computational Complexity*, 2003.

[26] C. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computing and Sys. Sci.*, 43:425–440, 1991.

[27] S. Sahni and T. Gonzales. P-complete approximation problems. *Journal of the ACM*, 23:555–565, 1976.

[28] A. Samorodnitsky and L. Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *Proc. 32nd Ann. ACM Symp. on the Theory of Computing*, pages 191–199, 2000.

[29] W. Sheppard. On the application of the theory of error to cases of normal distribution and normal correlation. *Phil. Trans. Royal Soc. London*, 192:101–168, 1899.

[30] M. Talagrand. How much are increasing sets positively correlated? *Combinatorica*, 16(2):243–258, 1996.

[31] L. Trevisan, G. Sorkin, M. Sudan, and D. Williamson. Gadgets, approximation, and linear programming. *SIAM Journal on Computing*, 29(6):2074–2097, 2000.

[32] U. Zwick. Outward rotations: a tool for rounding solutions of semidefinite programming relaxations, with applications to MAX-CUT and other problems. In *Proc. 31st Ann. ACM Symp. on the Theory of Computing*, pages 679–687, 1999.