Constraint Satisfaction over a Non-Boolean Domain: Approximation Algorithms and Unique-Games Hardness

Venkatesan Guruswami* and Prasad Raghavendra**

Department of Computer Science & Engineering,
University of Washington,
Seattle, WA
{venkat,prasad}@cs.washington.edu

Abstract. We study the approximability of the MAX k-CSP problem over non-boolean domains, more specifically over $\{0, 1, \ldots, q-1\}$ for some integer q. We extend the techniques of Samorodnitsky and Trevisan [19] to obtain a UGC hardness result when q is a prime. More precisely, assuming the Unique Games Conjecture, we show that it is NP-hard to approximate the problem to a ratio greater than q^2k/q^k. Independent of this work, Austrin and Mossel [2] obtain a more general UGC hardness result using entirely different techniques.

We also obtain an approximation algorithm that achieves a ratio of $C(q) \cdot k/q^k$ for some constant $C(q)$ depending only on q, via a subroutine for approximating the value of a semidefinite quadratic form when the variables take values on the corners of the q-dimensional simplex. This generalizes an algorithm of Nesterov [16] for the ± 1-valued variables. It has been pointed out to us [15] that a similar approximation ratio can be obtained by reducing the non-boolean case to a boolean CSP.

1 Introduction

Constraint Satisfaction Problems (CSP) capture a large variety of combinatorial optimization problems that arise in practice. In the MAX k-CSP problem, the input consists of a set of variables taking values over a domain(say $\{0, 1\}$), and a set of constraints with each acting on k of the variables. The objective is to find an assignment of values to the variables that maximizes the number of constraints satisfied. Several classic optimization problems like 3-SAT, Max Cut fall in to the general framework of CSPs. For most CSPs of interest, the problem of finding the optimal assignment turns out to be NP-hard. To cope with this intractability, the focus shifts to approximation algorithms with provable guarantees. Specifically, an algorithm $\mathcal{A}$ is said to yield an α approximation to a CSP, if on every instance Γ of the CSP, the algorithm outputs an assignment that satisfies at least α times as many constraints as the optimal assignment.

* Work done while on leave at School of Mathematics, Institute for Advanced Study, Princeton, NJ. Research supported in part by a Packard Fellowship and NSF grant CCF-0324906 to the IAS.

** Research supported by NSF CCF-0343672.

A. Goel et al. (Eds.): APPROX and RANDOM 2008, LNCS 5171, pp. 77–90, 2008.

Apart from its natural appeal, the study of the MAX k-CSP problem is interesting for yet another reason. The best approximation ratio achievable for MAX k-CSP equals the optimal soundness of a PCP verifier making at most k queries. In fact, inapproximability results for MAX k-CSP have often been accompanied by corresponding developments in analysis of linearity testing.

Over the boolean domain, the problem of MAX k-CSP has been studied extensively. For a boolean predicate $P : \{0,1\}^k \rightarrow \{0,1\}$, the MAX k-CSP (P) problem is the special case of MAX k-CSP where all the constraints are of the form $P(l_1, l_2, \ldots, l_k)$ with each literal l_i being either a variable or its negation. For many natural boolean predicates P, approximation algorithms and matching NP-hardness results are known for MAX k-CSP (P)[11]. For the general MAX k-CSP problem over boolean domain, the best known algorithm yields a ratio of $\Omega(\frac{k}{2^k})$ [3], while any ratio better than $2^{\sqrt{2k}}/2^k$ is known to be NP-hard to achieve [5]. Further if one assumes the Unique Games Conjecture, then it is NP-hard to approximate MAX k-CSP problem to a factor better than $\frac{2k}{2^k}$ [19].

In this work, we study the approximability of the MAX k-CSP problem over non-boolean domains, more specifically over $\{0, 1, \ldots, q-1\}$ for some integer q, obtaining both algorithmic and hardness results (under the UGC) with almost matching approximation factors.

On the hardness side, we extend the techniques of [19] to obtain a UGC hardness result when q is a prime. More precisely, assuming the Unique Games Conjecture, we show that it is NP-hard to approximate the problem to a ratio greater than q^2k/q^k. Except for constant factors depending on q, the algorithm and the UGC hardness result have the same dependence on of the arity k. Independent of this work, Austrin and Mossel [2] obtain a more general UGC hardness result using entirely different techniques. Technically, our proof extends the Gowers Uniformity based approach of Samorodnitsky and Trevisan [19] to correlations on q-ary cubes instead of the binary cube. This is related to the detection of multidimensional arithmetic progressions by a Gowers norm of appropriately large degree. Along the way, we also make a simplification to [19] and avoid the need to obtain a large *cross-influence* between two functions in a collection with a substantial Uniformity norm; instead our proof works based on large influence of just one function in the collection.

On the algorithmic side, we obtain a approximation algorithm that achieves a ratio of $C(q) \cdot k/q^k$ with $C(q) = \frac{1}{2\pi e q(q-1)^6}$. As a subroutine, we design an algorithm for maximizing a positive definite quadratic form with variables forced to take values on the corners of the q-dimensional simplex. This is a generalization of an algorithm of Nesterov [16] for maximizing positive definite quadratic form with variables forced to take $\{-1, 1\}$ values. Independent of this work, Makarychev and Makarychev [15] brought to our notice a reduction from non-boolean CSPs to the boolean case, which in conjunction with the CMM algorithm [3] yields a better approximation ratio for the MAX k-CSP problem. Using the reduction, one can deduce a $q^2(1 + o(1))k/q^k$ factor UG hardness for MAX k-CSP for arbitrary positive integers q, starting from our UG hardness result for primes q.

1.1 Related Work

The simplest algorithm for MAX k-CSP over boolean domain is to output a random assignment to the variables, thus achieving an approximation ratio of $\frac{1}{2^k}$. The first improvement over this trivial algorithm, a ratio of $\frac{2}{2^k}$ was obtained by Trevisan [20]. Hast [9] proposed an approximation algorithm with a ratio of $\Omega(\frac{k}{\log k 2^k})$, which was later improved to the current best known algorithm achieving an approximation factor of $\Omega(\frac{k}{2^k})$ [3].

On the hardness side, MAX k-CSP over the boolean domain was shown to be NP-hard to approximate to a ratio greater than $\Omega(2^{2\sqrt{k}}/2^k)$ by Samorodnitsky and Trevisan [18]. The result involved an analysis of a graph-linearity test which was simplified subsequently by Håstad and Wigderson [13]. Later, using the machinery of multi-layered PCP developed in [4], the inapproximability factor was improved to $O(2^{\sqrt{2k}}/2^k)$ in [5].

A predicate P is *approximation resistant* if the best optimal approximation ratio for MAX k-CSP (P) is given by the random assignment. While no predicate over 2 variables is approximation resistant, a predicate over 3 variables is approximation resistant if and only if it is implied by the XOR of 3 variables [11,21]. Almost all predicates on 4 variables were classified with respect to approximation resistance in [10].

In recent years, several inapproximability results for MAX k-CSP problems were obtained assuming the Unique Games Conjecture. Firstly, a tight inapproximability of $\Theta\left(\frac{k}{2^k}\right)$ was shown in [19]. The proof relies on the analysis of a hypergraph linearity test using the Gowers uniformity norms. Hastad showed that if UGC is true, then as k increases, nearly every predicate P on k variables is *approximation resistant* [12].

More recently, optimal inapproximability results have been shown for large classes of CSPs assuming the Unique Games Conjecture. Under an additional conjecture, optimal inapproximability results were obtained in [1] for all boolean predicates over 2 variables. Subsequently, it was shown in [17] that for every CSP over an arbitrary finite domain, the best possible approximation ratio is equal to the integrality gap of a well known Semidefinite program. Further the same work also obtains an algorithm that achieves the best possible approximation ratio assuming UGC. Although the results of [17] apply to non-boolean domains, they do not determine the value of the approximation factor explicitly, but only show that it is equal to the integrality gap of an SDP. Further the algorithm proposed in [17] does not yield any approximation guarantee for MAX k-CSP unconditionally. Thus neither the inapproximability nor the algorithmic results of this work are subsumed by [17].

Austrin and Mossel [2] obtain a sufficient condition for a predicate P to be approximation resistant. Through this sufficiency condition, they obtain strong UGC hardness results for MAX k-CSP problem over the domain $\{1, \ldots, q\}$ for arbitrary k and q. For the case when q is a prime power, their results imply a UGC hardness of $kq(q-1)/q^k$. The hardness results in this work and [2] were obtained independently and use entirely different techniques.

1.2 Organization of the Paper

We begin with background on the Unique Games conjecture, Gowers norm, and influence of variables in Section 2. In Section 3, we present a linearity test that forms the core of the UGC based hardness reduction. We prove our inapproximability result (for the case when q is a prime) by a reduction from Unique Games in Section 4. The proof uses a technical step bounding a certain expectation by an appropriate Gowers norm; this step is proved in Section 5. Finally, we state the algorithmic result in Section 6, deferring the details to the full version [6].

2 Preliminaries

In this section, we will set up notation, and review the notions of Gower's uniformity, influences, noise operators and the Unique games conjecture. Henceforth, for a positive integer n, we use the notation $[n]$ for the ring $\mathbb{Z}/(n) = \{0, 1, \dots, n-1\}$.

2.1 Unique Games Conjecture

Definition 1. *An instance of Unique Games represented as $\Gamma = (\mathcal{X} \cup \mathcal{Y}, E, \Pi, \langle R \rangle)$, consists of a bipartite graph over node sets $\mathcal{X}$,$\mathcal{Y}$ with the edges E between them. Also part of the instance is a set of labels $\langle R \rangle = \{1, \dots, R\}$, and a set of permutations $\pi_{vw} : \langle R \rangle \to \langle R \rangle$ for each edge $e = (v, w) \in E$. An assignment A of labels to vertices is said to satisfy an edge $e = (v, w)$, if $\pi_{vw}(A(v)) = A(w)$. The objective is to find an assignment A of labels that satisfies the maximum number of edges.*

For sake of convenience, we shall use the following stronger version of Unique Games Conjecture which is equivalent to the original conjecture [14].

Conjecture 1. For all constants $\delta > 0$, there exists large enough constant R such that given a bipartite unique games instance $\Gamma = (\mathcal{X} \cup \mathcal{Y}, E, \Pi = \{\pi_e : \langle R \rangle \to \langle R \rangle \ : \ e \in E\}, \langle R \rangle)$ with number of labels R, it is NP-hard to distinguish between the following two cases:

– $(1-\delta)$-satisfiable instances: There exists an assignment A of labels such that for $1-\delta$ fraction of vertices $v \in \mathcal{X}$, *all* the edges (v, w) are satisfied.
– Instances that are not δ-satisfiable: No assignment satisfies more than a δ-fraction of the edges E.

2.2 Gowers Uniformity Norm and Influence of Variables

We now recall the definition of the Gowers uniformity norm. For an integer $d \geqslant 1$ and a complex-valued function $f : G \to \mathbb{C}$ defined on an abelian group G (whose group operation we denote by $+$), the d'th uniformity norm $U_d(f)$ is defined as

$$U^d(f) := \mathop{\mathbb{E}}_{x, y_1, y_2, \dots, y_d} \left[\prod_{\substack{S \subseteq \{1,2,\dots,d\} \\ |S| \text{ even}}} f\left(x + \sum_{i \in S} y_i\right) \prod_{\substack{S \subseteq \{1,2,\dots,d\} \\ |S| \text{ odd}}} \overline{f\left(x + \sum_{i \in S} y_i\right)} \right] . \tag{1}$$

where the expectation is taken over uniform and independent choices of $x, y_0, \ldots, y_{d-1}$ from the group G. Note that $U^1(f) = \left(\mathbb{E}_x[f(x)]\right)^2$.

We will be interested in the case when the group G is $[q]^R$ for positive integers q, R, with group addition being coordinate-wise addition modulo q. G is also closed under coordinate-wise multiplication modulo q by scalars in $[q]$, and thus has a $[q]$-module structure. For technical reasons, we will restrict attention to the case when q is prime and thus our groups will be vector spaces over the field $\mathbb{F}_q$ of q elements. For a vector $\mathbf{a} \in [q]^k$, we denote by $a_1, a_2, \ldots, a_k$ its k coordinates. We will use $\mathbf{1}, \mathbf{0}$ to denote the all 1's and all 0's vectors respectively (the dimension will be clear from the context). Further denote by $\mathbf{e}_i$ the i^{th} basis vector with 1 in the i^{th} coordinate and 0 in the remaining coordinates. As we shall mainly be interested in functions over $[q]^R$ for a prime q, we make our further definitions in this setting. Firstly, every function $f : [q]^R \to \mathbb{C}$ has a Fourier expansion given by $f(x) = \sum_{\alpha \in [q]^R} \hat{f}_\alpha \chi_\alpha(x)$ where $\hat{f}_\alpha = \mathbb{E}_{x \in [q]^R}[f(x)\chi_\alpha(x)]$ and $\chi_\alpha(x) = \prod_{i=1}^R \omega^{\alpha_i x_i}$ for a q^{th} root of unity ω.

The central lemma in the hardness reduction relates a large Gowers norm for a function f, to the existence of an influential coordinate. Towards this, we define influence of a coordinate for a function over $[q]^R$.

Definition 2. *For a function $f : [q]^R \to \mathbb{C}$ define the influence of the i^{th} coordinate as follows:*

$$\mathrm{Inf}_i(f) = \mathbb{E}_x[\mathbf{Var}_{x_i}[f]] \ .$$

The following well known result relates influences to the Fourier spectrum of the function.

Fact 1. *For a function $f : [q]^R \to \mathbb{C}$ and a coordinate $i \in \{1, 2, \ldots, R\}$,*

$$\mathrm{Inf}_i(f) = \sum_{\alpha_i \neq 0, \alpha \in [q]^R} |\hat{f}_\alpha|^2 \ .$$

The following lemma is a restatement of Theorem 12 in [19].

Lemma 1. *There exists an absolute constant C such that, if $f : [q]^m \to \mathbb{C}$ is a function satisfying $|f(x)| \leqslant 1$ for every x then for every $d \geqslant 1$,*

$$U^d(f) \leqslant U^1(f) + 2^{Cd} \max_i \mathrm{Inf}_i(f)$$

2.3 Noise Operator

Like many other UGC hardness results, one of the crucial ingredients of our reduction will be a noise operator on functions over $[q]^R$. We define the noise operator $T_{1-\varepsilon}$ formally below.

Definition 3. *For $0 \leqslant \varepsilon \leqslant 1$, define the operator $T_{1-\varepsilon}$ on functions $f : [q]^R \to \mathbb{C}$ as:*

$$T_{1-\varepsilon} f(\mathbf{x}) = \mathbb{E}_\eta[f(\mathbf{x} + \eta)]$$

where each coordinate η_i of η is 0 with probability $1-\varepsilon$ and a random element from $[q]$ with probability ε. The Fourier expansion of $T_{1-\varepsilon}f$ is given by

$$T_{1-\varepsilon}f(\mathbf{x}) = \sum_{\alpha\in[q]^R} (1-\varepsilon)^{|\alpha|}\hat{f}_\alpha \chi_\alpha(x)$$

Here $|\alpha|$ denotes the number of non-zero coordinates of α. Due to space constraints, we defer the proof of the following lemma(see [6]).

Lemma 2. *If a function $f : [q]^R \to \mathbb{C}$ satisfies $|f(x)| \leqslant 1$ for all x, and $g = T_{1-\varepsilon}f$ then $\sum_{i=1}^{R} \mathrm{Inf}_i(g) \leqslant \frac{1}{2e\ln 1/(1-\varepsilon)}$*

3 Linearity Tests and MAX k-CSP Hardness

The best approximation ratio possible for MAX k-CSP is identical to the best soundness of a PCP verifier for NP that makes k queries. This follows easily by associating the proof locations to CSP variables, and the tests of the verifier to k-ary constraints on the locations. In this light, it is natural that the hardness results of [18,5,19] are all associated with a linearity test with a strong soundness. The hardness result in this work is obtained by extending the techniques of [19] from binary to q-ary domains. In this section, we describe the test of [19] and outline the extension to it.

For the sake of simplicity, let us consider the case when $k = 2^d - 1$ for some d. In [19], the authors propose the following linearity test for functions $F : \{0,1\}^n \to \{0,1\}$.

> Complete Hypergraph Test (F, d)
>
> – Pick $x_1, x_2, \ldots, x_d \in \{0,1\}^n$ uniformly at random.
> – Accept if for each $S \subseteq [r]$, $F(\sum_{i\in S} x_i) = \sum_{i\in S} F(x_i)$.

The test reads the value of the function F at $k = 2^d - 1$ points of a random subspace(spanned by $x_1, \ldots, x_d$) and checks that F agrees with a linear function on the subspace. Note that a random function F would pass the test with probability $2^d/2^k$, since there are 2^d different satisfying assignments to the k binary values queried by the verifier. The following result is a special case of a more general result by Samorodnitsky and Trevisan [19].

Theorem 1. *[19] If a function $F : \{0,1\}^n \to \{0,1\}$ passes the Complete Hypergraph Test with probability greater than $2^d/2^k + \gamma$, then the function $f(x) = (-1)^{F(x)}$ has a large d^{th} Gowers norm. Formally, $U^d(f) \geqslant C(\gamma, k)$ for some fixed function C of γ, k.*

Towards extending the result to the domain $[q]$, we propose a different linearity test. Again for convenience, let us assume $k = q^d$ for some d. Given a function $F : [q]^n \to [q]$, the test proceeds as follows:

Affine Subspace Test (F, d)

- Pick $\mathbf{x}, \mathbf{y_1}, \mathbf{y_2}, \ldots, \mathbf{y_d} \in [q]^n$ uniformly at random.
- Accept if for each $\mathbf{a} \subseteq [q]^d$,

$$F\Big(\mathbf{x} + \sum_{i=1}^{d} a_i \mathbf{y_i}\Big) = \Big(1 - \sum_{i=1}^{d} a_i\Big) F(\mathbf{x}) + \sum_{i=1}^{d} a_i F\Big(\mathbf{x} + \mathbf{y_i}\Big)$$

Essentially, the test queries the values along a randomly chosen affine subspace, and tests if the function F agrees with an affine function on the subspace. Let ω denote a q'th root of unity. From Theorem 4 presented in Section 5, the following result can be shown:

Theorem 2. *If a function $F : [q]^n \to [q]$ passes the Affine Subspace Test with probability greater than $q^{d+1}/q^k + \gamma$, then for some q'th root of unity $\omega \neq 1$, the function $f(x) = \omega^{F(x)}$ has a large dq'th Gowers norm . Formally, $U^{dq}(f) \geqslant C(\gamma, k)$ for some fixed function C of γ, k.*

The above result follows easily from Theorem 4 using techniques of [19], and the proof is ommited here. The Affine Subspace Test forms the core of the UGC based hardness reduction presented in Section 4.

4 Hardness Reduction from Unique Games

In this section, we will prove a hardness result for approximating MAX k-CSP over a domain of size q when q is prime for every $k \geqslant 2$. Let d be such that $q^{d-1} + 1 \leqslant k \leqslant q^d$. Let us consider the elements of $[q]$ to have a natural order defined by $0 < 1 < \ldots < q-1$. This extends to a lexicographic ordering on vectors in $[q]^d$. Denote by $[q]^d_{<k}$ the set consisting of the k lexicographically smallest vectors in $[q]^d$. We shall identify the set $\{1, \ldots, k\}$ with set of vectors in $[q]^d_{<k}$. Specifically, we shall use $\{1, \ldots, k\}$ and vectors in $[q]^d_{<k}$ interchangeably as indices to the same set of variables. For a vector $\mathbf{x} \in [q]^R$ and a permutation π of $\{1, \ldots, R\}$, define $\pi(x) \in [q]^R$ defined by $(\pi(x))_i = x_{\pi(i)}$.

Let $\Gamma = (\mathcal{X} \cup \mathcal{Y}, E, \Pi = \{\pi_e : \langle R \rangle \to \langle R \rangle | e \in E\}, \langle R \rangle)$ be a bipartite unique games instance. Towards constructing a k-CSP instance Λ from Γ, we shall introduce a long code for each vertex in $\mathcal{Y}$. Specifically, the set of variables for the k-CSP Λ is indexed by $\mathcal{Y} \times [q]^R$. Thus a solution to Λ consists of a set of functions $F_w : [q]^R \to [q]$, one for each $w \in \mathcal{Y}$.

Similar to several other long code based hardness results, we shall assume that the long codes are *folded*. More precisely, we shall use *folding* to force the functions F_w to satisfy $F_w(\mathbf{x} + \mathbf{1}) = F(\mathbf{x}) + 1$ for all $\mathbf{x} \in [q]^R$. The k-ary constraints in the instance Λ are specified by the following verifier. The verifier uses an additional parameter ε that governs the level of noise in the noise operator.

- Pick a random vertex $v \in \mathcal{X}$. Pick k vertices $\{w_{\mathbf{a}} | \mathbf{a} \in [q]^d_{<k}\}$ from $N(v) \subset \mathcal{Y}$ uniformly at random independently. Let $\pi_{\mathbf{a}}$ denote the permutation on the edge $(v, w_{\mathbf{a}})$.
- Sample $\mathbf{x}, \mathbf{y_1}, \mathbf{y_2}, \ldots, \mathbf{y_d} \in [q]^R$ uniformly at random. Sample vectors $\eta_{\mathbf{a}} \in [q]^R$ for each $\mathbf{a} \in [q]^d_{<k}$ from the following distribution: With probability $1-\varepsilon$, $(\eta_{\mathbf{a}})_j = 0$ and with the remaining probability, $(\eta_{\mathbf{a}})_j$ is a uniformly random element from $[q]$.
- Query $F_{w_{\mathbf{a}}}\Big(\pi_{\mathbf{a}}(\mathbf{x} + \sum_j a_j \mathbf{y_j} + \eta_{\mathbf{a}})\Big)$ for each $\mathbf{a} \in [q]^d_{<k}$. Accept if the following equality holds for each $\mathbf{a} \in [q]^d_{<k}$.

$$F_{w_{\mathbf{a}}}\Big(\pi_{\mathbf{a}}(\mathbf{x} + \sum_{j=1}^{d} a_j \mathbf{y_j} + \eta_{\mathbf{a}})\Big) = \Big(1 - \sum_{j=1}^{d} a_j\Big) F_{w_{\mathbf{0}}}\Big(\pi_{\mathbf{0}}(\mathbf{x} + \eta_{\mathbf{0}})\Big) + \sum_{j=1}^{d} a_j F_{w_{\mathbf{e_j}}}\Big(\pi_{\mathbf{e_j}}(\mathbf{x} + \mathbf{y_j} + \eta_{\mathbf{e_j}})\Big)$$

Theorem 3. *For all primes q, positive integers d, k satisfying $q^{d-1} < k \leqslant q^d$, and every $\gamma > 0$, there exists small enough $\delta, \varepsilon > 0$ such that*

- Completeness: *If Γ is a $(1-\delta)$-satisfiable instance of Unique Games, then there is an assignment to Λ that satisfies the verifier's tests with probability at least $(1-\gamma)$*
- Soundness: *If Γ is not δ-satisfiable, then no assignment to Λ satisfies the verifier's tests with probability more than $\frac{q^{d+1}}{q^k} + \gamma$.*

Proof. We begin with the completeness claim, which is straightforward.

Completeness. There exists labelings to the Unique Game instance Γ such that for $1-\delta$ fraction of the vertices $v \in \mathcal{X}$ all the edges (v, w) are satisfied. Let $A : \mathcal{X} \cup \mathcal{Y} \to \langle R \rangle$ denote one such labelling. Define an assignment to the k-CSP instance by $F_w(\mathbf{x}) = x_{A(w)}$ for all $w \in \mathcal{Y}$.

With probability at least $(1-\delta)$, the verifier picks a vertex $v \in \mathcal{X}$ such that the assignment A satisfies all the edges $(v, w_{\mathbf{a}})$. In this case for each $\mathbf{a}$, $\pi_{\mathbf{a}}(A(v)) = A(w_{\mathbf{a}})$. Let us denote $A(v) = l$. By definition of the functions F_w, we get $F_{w_{\mathbf{a}}}(\pi_{\mathbf{a}}(x)) = (\pi_{\mathbf{a}}(x))_{A(w_{\mathbf{a}})} = x_{\pi_{\mathbf{a}}^{-1}(A(w_{\mathbf{a}}))} = x_l$ for all $x \in [q]^R$. With probability at least $(1-\varepsilon)^k$, each of the vectors $\eta_{\mathbf{a}}$ have their l^{th} component equal to zero, i.e $(\eta_{\mathbf{a}})_l = 0$. In this case, it is easy to check that all the constraints are satisfied. In conclusion, the verifier accepts the assignment with probability at least $(1-\delta)(1-\varepsilon)^k$. For small enough δ, ε, this quantity is at least $(1-\gamma)$.

Soundness. Suppose there is an assignment given by functions F_w for $w \in \mathcal{Y}$ that the verifier accepts with probability greater than $\frac{q^{d+1}}{q^k} + \gamma$.

Let $z_1, z_2, \ldots, z_k$ be random variables denoting the k values read by the verifier. Thus $z_1, \ldots, z_k$ take values in $[q]$. Let $P : [q]^k \to \{0, 1\}$ denote the predicate on k variables that represents the acceptance criterion of the verifier. Essentially, the value of the predicate $P(z_1, \ldots, z_k)$ is 1 if and only if $z_1, \ldots, z_k$ values are consistent with some affine function. By definition,

$$\Pr[\text{ Verifier Accepts }] = \mathop{\mathbb{E}}_{v \in \mathcal{X}} \mathop{\mathbb{E}}_{w_{\mathbf{a}} \in N(v)} \mathop{\mathbb{E}}_{\mathbf{x}, \mathbf{y_1}, \ldots, \mathbf{y_d}} \mathop{\mathbb{E}}_{\eta_{\mathbf{a}}} \Big[P(z_1, \ldots, z_k) \Big] \geqslant \frac{q^{d+1}}{q^k} + \gamma$$

Let ω denote a q^{th} root of unity. The Fourier expansion of the function $P : [q]^k \to \mathbb{C}$ is given by $P(z_1, \ldots, z_k) = \sum_{\alpha \in [q]^k} \hat{P}_\alpha \chi_\alpha(z_1, \ldots, z_k)$ where $\chi_\alpha(z_1, \ldots, z_k) = \prod_{i=1}^k \omega^{\alpha_i z_i}$ and $\hat{P}_\alpha = \mathop{\mathbb{E}}_{z_1, \ldots, z_k} [P(z_1, \ldots, z_k)\chi_\alpha(z_1, \ldots, z_k)]$. Notice that for $\alpha = \mathbf{0}$, we get $\chi_\alpha(z_1, \ldots, z_k) = 1$. Further,

$$\hat{P}_{\mathbf{0}} = \Pr[\text{ random assignment to } z_1, z_2, \ldots, z_k \text{ satisfies } P] = \frac{q^{d+1}}{q^k}$$

Substituting the Fourier expansion of P, we get

$$\Pr[\text{ Verifier Accepts }] = \frac{q^{d+1}}{q^k} + \sum_{\alpha \neq \mathbf{0}} \hat{P}_\alpha \mathop{\mathbb{E}}_{v \in \mathcal{X}} \mathop{\mathbb{E}}_{w_{\mathbf{a}} \in N(v)} \mathop{\mathbb{E}}_{\mathbf{x}, \mathbf{y_1}, \ldots, \mathbf{y_d}} \mathop{\mathbb{E}}_{\eta_{\mathbf{a}}} \Big[\chi_\alpha(z_1, \ldots, z_k) \Big]$$

Recall that the probability of acceptance is greater than $\frac{q^{d+1}}{q^k} + \gamma$. Further $|\hat{P}_\alpha| \leqslant 1$ for all $\alpha \in [q]^k$. Thus there exists $\alpha \neq 0$ such that,

$$\Big| \mathop{\mathbb{E}}_{v \in \mathcal{X}} \mathop{\mathbb{E}}_{w_{\mathbf{a}} \in N(v)} \mathop{\mathbb{E}}_{\mathbf{x}, \mathbf{y_1}, \ldots, \mathbf{y_d}} \mathop{\mathbb{E}}_{\eta_{\mathbf{a}}} \Big[\chi_\alpha(z_1, \ldots, z_k) \Big] \Big| \geqslant \frac{\gamma}{q^k}$$

For each $w \in \mathcal{Y}, t \in [q]$, define the function $f_w^{(t)} : [q]^d \to \mathbb{C}$ as $f_w^{(t)}(x) = \omega^{tF_w(x)}$. For convenience we shall index the vector α with the set $[q]^d_{<k}$ instead of $\{1, \ldots, k\}$. In this notation,

$$\Big| \mathop{\mathbb{E}}_{v \in \mathcal{X}} \mathop{\mathbb{E}}_{w_{\mathbf{a}} \in N(v)} \mathop{\mathbb{E}}_{\mathbf{x}, \mathbf{y_1}, \ldots, \mathbf{y_d}} \mathop{\mathbb{E}}_{\eta_{\mathbf{a}}} \Big[\prod_{\mathbf{a} \in [q]^d_{<k}} f_{w_{\mathbf{a}}}^{(\alpha_{\mathbf{a}})} \big(\pi_{\mathbf{a}}(\mathbf{x} + \sum_{i=1}^d a_i \mathbf{y_i} + \eta_{\mathbf{a}}) \big) \Big] \Big| \geqslant \frac{\gamma}{q^k}$$

Let $g_w^{(t)} : [q]^d \to \mathbb{C}$ denote the *smoothened* version of function $f_w^{(t)}$. Specifically, let $g_w^{(t)}(x) = T_{1-\varepsilon} f_w^{(t)}(x) = \mathbf{E}_\eta[f_w^{(t)}(x + \eta)]$ where η is generated from ε-noise distribution. Since each $\eta_{\mathbf{a}}$ is independently chosen, we can rewrite the above expression,

$$\Big| \mathop{\mathbb{E}}_{v \in \mathcal{X}} \mathop{\mathbb{E}}_{w_{\mathbf{a}} \in N(v)} \mathop{\mathbb{E}}_{\mathbf{x}, \mathbf{y_1}, \ldots, \mathbf{y_d}} \Big[\prod_{\mathbf{a} \in [q]^d_{<k}} g_{w_{\mathbf{a}}}^{(\alpha_{\mathbf{a}})} \big(\pi_{\mathbf{a}}(\mathbf{x} + \sum_{i=1}^d a_i \mathbf{y_i}) \big) \Big] \Big| \geqslant \frac{\gamma}{q^k} \ .$$

For each $v \in \mathcal{X}, t \in [q]$, define the function $g_v^{(t)} : [q]^d \to \mathbb{C}$ as $g_v^{(t)}(x) = \mathbf{E}_{w \in N(v)}[g_w^{(t)}(\pi_{vw}(x))]$. As the vertices $w_{\mathbf{a}}$ are chosen independent of each other,

$$\Big| \mathop{\mathbb{E}}_{v \in \mathcal{X}} \mathop{\mathbb{E}}_{\mathbf{x}, \mathbf{y_1}, \ldots, \mathbf{y_d}} \Big[\prod_{\mathbf{a} \in [q]^d_{<k}} g_v^{(\alpha_{\mathbf{a}})} \big(\mathbf{x} + \sum_{i=1}^d a_i \mathbf{y_i} \big) \Big] \Big| \geqslant \frac{\gamma}{q^k} \ .$$

As $\alpha \neq 0$, there exists an index $\mathbf{b} \in [q]^d_{<k}$ such that $\alpha_{\mathbf{b}} \neq 0$. For convenience let us denote $c = \alpha_{\mathbf{b}}$. Define $\kappa = 2^{-Cdq}\left(\frac{\gamma}{2q^k}\right)^{2^{dq}}$ where C is the absolute constant defined in Lemma 1.

For each $v \in \mathcal{X}$, define the set of labels $L(v) = \{i \in \langle R\rangle \ : \ \mathrm{Inf}_i(g_v^c) \geqslant \kappa\}$. Similarly for each $w \in \mathcal{Y}$, let $L(w) = \{i \in \langle R\rangle \ : \ \mathrm{Inf}_i(g_w^c) \geqslant \kappa/2\}$. Obtain a labelling A to the Unique Games instance Γ as follows : For each vertex $u \in \mathcal{X} \cup \mathcal{Y}$, if $L(u) \neq \phi$ then assign a randomly chosen label from $L(u)$, else assign a uniformly random label from $\langle R\rangle$.

The functions $g_w^{(c)}$ are given by $g_w^{(c)} = T_{1-\varepsilon} f_w^{(c)}$ where $f_w^{(c)}$ is bounded in absolute value by 1. By Lemma 2, therefore, the sum of its influences is bounded by $\frac{1}{e \ln 1/(1-\varepsilon)}$. Consequently, for all $w \in \mathcal{Y}$ the size of the label set $L(w)$ is bounded by $\frac{2}{\kappa e \ln 1/(1-\varepsilon)}$. Applying a similar argument to $v \in \mathcal{X}$, $|L(v)| \leqslant \frac{1}{\kappa e \ln 1/(1-\varepsilon)}$.

For at least $\gamma/2q^k$ fraction of vertices $v \in \mathcal{X}$ we have,

$$\Big| \mathop{\mathbb{E}}_{\mathbf{x},\mathbf{y_1},\ldots,\mathbf{y_d}} \Big[\prod_{\mathbf{a} \in [q]^d_{<k}} g_v^{(\alpha_{\mathbf{a}})}\big(\mathbf{x} + \sum_{i=1}^{d} a_i \mathbf{y_i}\big)\Big]\Big| \geqslant \frac{\gamma}{2q^k}$$

We shall refer to these vertices as *good* vertices. Fix a *good* vertex v.

Observe that for each $u \in \mathcal{X} \cup \mathcal{Y}$ the functions $g_u^{(t)}$ satisfy $|g_u^{(t)}(x)| \leqslant 1$ for all x. Now we shall apply Theorem 4 to conclude that the functions $g_v^{(t)}$ have a large Gowers norm. Specifically, consider the collection of functions given by $f_{\mathbf{a}} = g_v^{(\alpha_{\mathbf{a}})}$ for $\mathbf{a} \in [q]^d_{<k}$, and $f_{\mathbf{a}} = 1$ for all $\mathbf{a} \notin [q]^d_{<k}$. From Theorem 4, we get

$$\min_{\mathbf{a}} U^{dq}(g_v^{(\alpha_{\mathbf{a}})}) \geqslant \left(\frac{\gamma}{2q^k}\right)^{2^{dq}} .$$

In particular, this implies $U^{dq}(g_v^{(c)}) \geqslant \left(\frac{\gamma}{2q^k}\right)^{2^{dq}}$. Now we shall use Lemma 1 to conclude that the function g_v has influential coordinates. Towards this, observe that the functions $f_w^{(t)}$ satisfy $f_w^{(t)}(x+1) = f_w^{(t)}(x) \cdot \omega^t$ due to folding. Thus for all $t \neq 0$ and all $w \in \mathcal{Y}$, $\mathbf{E}_x[f_w^{(t)}(x)] = 0$. Specifically for $c \neq 0$,

$$U^1(g_v^{(c)}) = \Big(\mathop{\mathbb{E}}_{x}[g_v^{(c)}(x)]\Big)^2 = \Big(\mathop{\mathbb{E}}_{w \in N(v)} \mathop{\mathbb{E}}_{\eta} \mathop{\mathbb{E}}_{x}[f_w^{(c)}(x+\eta)]\Big)^2 = 0$$

Hence it follows from Lemma 1 that there exists influential coordinates i with $\mathrm{Inf}_i(g_v^{(c)}) \geqslant 2^{-Cdq}\left(\frac{\gamma}{2q^k}\right)^{2^{dq}} = \kappa$. In other words, $L(v)$ is non-empty. Observe that, due to convexity of influences,

$$\mathrm{Inf}_i(g_v^{(c)}) = \mathrm{Inf}_i(\mathop{\mathbb{E}}_{w \in N(v)}[g_w^{(c)}]) \leqslant \mathop{\mathbb{E}}_{w \in N(v)} \mathrm{Inf}_{\pi_{vw}(i)}([g_w^{(c)}(x)]) .$$

If the coordinate i has influence at least κ on $g_v^{(c)}$, then the coordinate $\pi_{vw}(i)$ has an influence of at least $\kappa/2$ for at least $\kappa/2$ fraction of neighbors $w \in N(v)$. The edge π_{vw} is satisfied if i is assigned to v, and $\pi_{wv}(i)$ is assigned to w. This event happens with probability at least $\frac{1}{|L(u)||L(v)|} \geqslant (e\kappa \ln 1/(1-\varepsilon))^2/2$ for at least $\kappa/2$ fraction of the neighbors $w \in N(v)$. As there are at least $(\gamma/2q^k)$ fraction of *good* vertices v, the assignment satisfies at least $(\gamma/2q^k)(e\kappa \ln 1/(1-\varepsilon))^2\kappa/4$ fraction of the unique games constraints. By choosing δ smaller than this fraction, the proof is complete.

Since each test performed by the verifier involve k variables, by the standard connection between hardness of MAX k-CSP and k-query PCP verifiers, we get the following hardness result conditioned on the UGC.

Corollary 1. *Assuming the Unique Games conjecture, for every prime q, it is NP-hard to approximate* MAX k-CSP *over domain size q within a factor that is greater than q^2k/q^k.*

Using the reduction of [15], the above UG hardness result can be extended from primes to arbitrary composite number q.

Corollary 2. *[15] Assuming the Unique Games conjecture, for every positive integer q, it is NP-hard to approximate* MAX k-CSP *over domain size q within a factor that is greater than $q^2k(1+o(1))/q^k$.*

5 Gowers Norm and Multidimensional Arithmetic Progressions

The following theorem forms a crucial ingredient in the soundness analysis in the proof of Theorem 3.

Theorem 4. *Let $q \geqslant 2$ be a prime and G be a $\mathbb{F}_q$-vector space. Then for all positive integers $\ell \leqslant q$ and d, and all collections $\{f_{\mathbf{a}} : G \to \mathbb{C}\}_{\mathbf{a}\in[\ell]^d}$ of ℓ^d functions satisfying $|f_{\mathbf{a}}(x)| \leqslant 1$ for every $x \in G$ and $\mathbf{a} \in [\ell]^d$, the following holds:*

$$\left| \mathop{\mathbb{E}}_{x,y_1,y_2,\ldots,y_d} \left[\prod_{\mathbf{a}\in[\ell]^d} f_{\mathbf{a}}(x + a_1y_1 + a_2y_2 + \cdots + a_dy_d) \right] \right| \leqslant \min_{\mathbf{a}\in[\ell]^d} \left(U^{d\ell}(f_{\mathbf{a}}) \right)^{1/2^{d\ell}} \tag{2}$$

The proof of the above theorem is via double induction on d, ℓ. We first prove the theorem for the one-dimensional case, i.e., $d = 1$ and every ℓ, $1 \leqslant \ell < q$ (Lemma 3). This will be done through induction on ℓ. We will then prove the result for arbitrary d by induction on d.

Remark 1. Green and Tao, in their work [8] on configurations in the primes, isolate and define a property of a system of linear forms that ensures that the degree t Gowers norm is sufficient to analyze patterns corresponding to those linear forms, and called this property *complexity* (see Definition 1.5 in [8]). Gowers and Wolf [7] later coined the term Cauchy-Schwartz (CS) complexity to refer to this notion of complexity. For example, the CS-complexity of the q linear forms $x, x+y, x+2y, \ldots, x+(q-1)y$ corresponding to a q-term arithmetic progression equals $q-2$, and the U^{q-1} norm suffices to analyze them. It can similarly be shown that the CS-complexity of the d-dimensional arithmetic progression (with q^d linear forms as in (2)) is at most $d(q-1)-1$. In our application, we need a "multi-function" version of these statements, since we have a different function $f_{\mathbf{a}}$ for each linear form $x+\mathbf{a}\cdot\mathbf{y}$. We therefore work out a self-contained proof of Theorem 4 in this setting.

Towards proving Theorem 4, we will need the following lemma whose proof is presented in the full version[6].

Lemma 3. *Let $q \geqslant 2$ be prime and ℓ, $1 \leqslant \ell \leqslant q$, be an integer, and G be a $\mathbb{F}_q$-vector space. Let $\{h_\alpha : G \to \mathbb{C}\}_{\alpha\in[\ell]}$ be a collection of ℓ functions such that $|h_\alpha(x)| \leqslant 1$ for all $\alpha \in [\ell]$ and $x \in G$. Then*

$$\left| \mathop{\mathbb{E}}_{x,y_1} \left[\prod_{\alpha\in[\ell]} h_\alpha(x+\alpha y_1) \right] \right| \leqslant \min_{\alpha\in[\ell]} \left(U^\ell(h_\alpha)\right)^{\frac{1}{2^\ell}} . \tag{3}$$

Proof of Theorem 4: Fix an arbitrary ℓ, $1 \leqslant \ell \leqslant q$. We will prove the result by induction on d. The base case $d=1$ is the content of Lemma 3, so it remains to consider the case $d>1$.

By a change of variables, it suffices to upper bound the LHS of (2) by $\left(U^{d\ell}(f_{(\ell-1)\mathbf{1}})\right)^{1/2^{d\ell}}$, and this is what we will prove. For $\alpha \in [\ell]$, and $y_2, y_3, \ldots, y_d \in G$, define the function

$$g_\alpha^{y_2,\ldots,y_d}(x) = \prod_{\mathbf{b}=(b_2,b_3,\ldots,b_d)\in[\ell]^{d-1}} f_{(\alpha,\mathbf{b})}(x+b_2y_2+\cdots+b_dy_d) . \tag{4}$$

The LHS of (2), raised to the power $2^{d\ell}$, equals

$$\begin{aligned}
\left| \mathop{\mathbb{E}}_{y_2,\ldots,y_d} \mathop{\mathbb{E}}_{x,y_1} \left[\prod_{\alpha\in[\ell]} g_\alpha^{y_2,\ldots,y_d}(x+\alpha y_1) \right] \right|^{2^{d\ell}} &\leqslant \left(\mathop{\mathbb{E}}_{y_2,\ldots,y_d} \left| \mathop{\mathbb{E}}_{x,y_1} \prod_{\alpha\in[\ell]} g_\alpha^{y_2,\ldots,y_d}(x+\alpha y_1) \right|^{2^\ell} \right)^{2^{(d-1)\ell}} \\
&\leqslant \left| \mathop{\mathbb{E}}_{y_2,\ldots,y_d} U^\ell(g_{\ell-1}^{y_2,\ldots,y_d}) \right|^{2^{(d-1)\ell}} \qquad \text{(using Lemma 3)} \\
&= \left| \mathop{\mathbb{E}}_{y_2,\ldots,y_d} \mathop{\mathbb{E}}_{x,z_1,\ldots,z_\ell} \left[\prod_{S\subseteq\{1,2,\ldots,\ell\}} g_{\ell-1}^{y_2,\ldots,y_d}\Big(x+\sum_{i\in S} z_i\Big) \right] \right|^{2^{(d-1)\ell}}
\end{aligned}$$

Defining the function

$$H_{\mathbf{b}}^{z_1,\ldots,z_\ell}(t) := \prod_{S\subseteq\{1,2,\ldots,\ell\}} f_{(\ell-1,\mathbf{b})}\Big(t+\sum_{i\in S} z_i\Big) \tag{5}$$

for every $\mathbf{b} \in [\ell]^{d-1}$ and $z_1, \ldots, z_\ell \in G$, the last expression equals

$$\left| \mathop{\mathbb{E}}_{z_1,\ldots,z_\ell} \mathop{\mathbb{E}}_{x,y_2,\ldots,y_d} \left[\prod_{\mathbf{b}=(b_2,\ldots,b_d)\in[\ell]^{d-1}} H_{\mathbf{b}}^{z_1,\ldots,z_\ell}\Big(x + b_2y_2 + \cdots + b_dy_d\Big)\right] \right|^{2^{(d-1)\ell}}$$

which is at most

$$\mathop{\mathbb{E}}_{z_1,\ldots,z_\ell} \left[\left| \mathop{\mathbb{E}}_{x,y_2,\ldots,y_d} \left[\prod_{\mathbf{b}=(b_2,\ldots,b_d)\in[\ell]^{d-1}} H_{\mathbf{b}}^{z_1,\ldots,z_\ell}\Big(x + b_2y_2 + \cdots + b_dy_d\Big)\right] \right|^{2^{(d-1)\ell}} \right]. \tag{6}$$

By the induction hypothesis, (6) is at most $\mathop{\mathbb{E}}_{z_1,\ldots,z_\ell}\left[U^{(d-1)\ell}\Big(H_{(\ell-1)\mathbf{1}}^{z_1,\ldots,z_\ell}\Big)\right]$. Recalling the definition of $H_{\mathbf{b}}^{z_1,\ldots,z_\ell}$ from (5), the above expectation equals

$$\mathop{\mathbb{E}}_{z_1,\ldots,z_\ell} \mathop{\mathbb{E}}_{\substack{x,\{z'_j\}\\ 1\leqslant j\leqslant(d-1)\ell}} \left[\prod_{\substack{S\subseteq\{1,2,\ldots,\ell\}\\ T\subseteq\{1,2,\ldots,(d-1)\ell\}}} f_{(\ell-1)\mathbf{1}}\Big(x + \sum_{i\in S} z_i + \sum_{j\in T} z'_j\Big)\right]$$

which clearly equals $U^{d\ell}(f_{(\ell-1)\mathbf{1}})$.

6 Approximation Algorithm for MAX k-CSP

On the algorithmic side, we show the following result:

Theorem 5. *There is a polynomial time algorithm that computes a* $\frac{1}{2\pi e q(q-1)^6} \cdot \frac{k}{q^k}$ *factor approximation for the* MAX k-CSP *problem over a domain of size* q.

The algorithm proceeds along the lines of [3], by formulating MAX k-CSP as a quadratic program, solving a SDP relaxation and rounding the resulting solution. The variables in the quadratic program are constrained to the vertices of the q-dimensional simplex. Hence, as a subroutine, we obtain an efficient procedure to optimize positive definite quadratic forms with the variables forced to take values on the q-dimensional simplex. Let Δ_q denote the q-dimensional simplex, and let $\mathsf{Vert}(\Delta_q)$ denote the vertices of the simplex. Formally,

Theorem 6. *Let* $A = (a_{ij}^{(k)(l)})$ *be a positive definite matrix where* $k, l \in [q]$ *and* $1 \leqslant i, j \leqslant n$. *For the quadratic program* Γ, *there exists an efficient algorithm that finds an assignment whose value is at least* $\frac{2}{\pi(q-1)^4}$ *of the optimum.*

QuadraticProgram Γ	
Maximize	$\sum_{ij} a_{ij}^{(k)(l)} x_i^{(k)} \cdot x_j^{(l)}$
Subject to	$\mathbf{x}_i = (x_i^{(0)}, x_i^{(1)}, \ldots, x_i^{(q-1)}) \in \mathsf{Vert}(\Delta_q)\ 1 \leqslant i \leqslant n$

The details of the algorithm are presented in the full version[6]. It has been pointed out to us that a $\Omega(q^2k/q^k)$-approximation for MAX k-CSP can be obtained by reducing from the non-boolean to the boolean case [15].

References

1. Austrin, P.: Towards sharp inapproximability for any 2-csp. In: FOCS: IEEE Symposium on Foundations of Computer Science, pp. 307–317 (2007)
2. Austrin, P., Mossel, E.: Approximation resistant predicates from pairwise independence. Electronic Colloqium on Computational Complexity, TR08-009 (2008)
3. Charikar, M., Makarychev, K., Makarychev, Y.: Near-optimal algorithms for maximum constraint satisfaction problems. In: Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 62–68 (2007)
4. Dinur, I., Guruswami, V., Khot, S., Regev, O.: A new multilayered PCP and the hardness of hypergraph vertex cover. SIAM J. Computing 34(5), 1129–1146 (2005)
5. Engebretsen, L., Holmerin, J.: More efficient queries in PCPs for NP and improved approximation hardness of maximum CSP. In: Diekert, V., Durand, B. (eds.) STACS 2005. LNCS, vol. 3404, pp. 194–205. Springer, Heidelberg (2005)
6. Guruswami, V., Raghavendra, P.: Constraint Satisfaction over the Non-Boolean Domain: Approximation algorithms and Unique Games hardness. ECCC: Electronic Colloqium on Computational Complexity, TR08-008 (2008)
7. Gowers, W.T., Wolf, J.: The true complexity of a system of linear equations. arXiv:math.NT/0711.0185 (2007)
8. Green, B., Tao, T.: Linear equations in primes. arXiv:math.NT/0606088v1 (2006)
9. Hast, G.: Approximating MAX kCSP – outperforming a random assignment with almost a linear factor. In: ICALP: Annual International Colloquium on Automata, Languages and Programming (2005)
10. Hast, G.: Beating a random assignment - approximating constraint satisfaction problems. Phd Thesis, Royal Institute of Technology (2005)
11. Håstad, J.: Some optimal inapproximability results. Journal of the ACM 48(4), 798–859 (2001)
12. Håstad, J.: On the approximation resistance of a random predicate. In: APPROX-RANDOM, pp. 149–163 (2007)
13. Håstad, J., Wigderson, A.: Simple analysis of graph tests for linearity and PCP. Random Struct. Algorithms 22(2), 139–160 (2003)
14. Khot, S., Regev, O.: Vertex cover might be hard to approximate to within /spl epsi/. In: Annual IEEE Conference on Computational Complexity (formerly Annual Conference on Structure in Complexity Theory), vol. 18 (2003)
15. Makarychev, K., Makarychev, Y.: Personal communication (2008)
16. Nesterov, Y.: Quality of semidefinite relaxation for nonconvex quadratic optimization. CORE Discussion Paper 9719 (1997)
17. Raghavendra, P.: Optimal algorithm and inapproximability results for every csp? In: STOC: ACM Symposium on Theory of Computing, pp. 245–254 (2008)
18. Samorodnitsky, A., Trevisan, L.: A PCP characterization of NP with optimal amortized query complexity. In: STOC: ACM Symposium on Theory of Computing, pp. 191–199 (2000)
19. Samorodnitsky, A., Trevisan, L.: Gowers uniformity, influence of variables, and PCPs. In: STOC: ACM Symposium on Theory of Computing (2006)
20. Trevisan, L.: Parallel approximation algorithms by positive linear programming. Algorithmica 21(1), 72–88 (1998)
21. Zwick, U.: Approximation algorithms for constraint satisfaction problems involving at most three variables per constraint. In: Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 201–210 (1998)