THEORY OF COMPUTING, Volume 12 (6), 2016, pp. 1–11
www.theoryofcomputing.org

NOTE

Simple Proof of Hardness of Feedback Vertex Set

Venkatesan Guruswami* Euiwoong Lee†

Received June 3, 2015; Revised February 7, 2016; Published August 2, 2016

Abstract: The Feedback Vertex Set problem (FVS), where the goal is to find a small subset of vertices that intersects every cycle in an input directed graph, is among the fundamental problems whose approximability is not well understood. One can efficiently find an $\widetilde{O}(\log n)$-factor approximation, and efficient constant-factor approximation is ruled out under the Unique Games Conjecture (UGC). We give a simpler proof that Feedback Vertex Set is hard to approximate within any constant factor, assuming UGC.

ACM Classification: F.2.2

AMS Classification: 68W25

Key words and phrases: inapproximability, UG-hardness, feedback vertex set

1 Introduction

Feedback Vertex Set (FVS) is a fundamental combinatorial optimization problem. Given a (directed) graph G, the problem asks to find a subset F of vertices[1] with the minimum cardinality that intersects every cycle in the graph (equivalently, the induced subgraph $G \setminus F$ is acyclic). One of Karp's 21 NP-complete problems, FVS has been a subject of active research for many years. Recent results on the

*Supported in part by CCF-1115525.

†Supported in part by CCF-1115525 and Samsung Scholarship.

[1]The related Feedback *Arc* Set problem asks for a subset of *edges* to intersect every cycle. This problem is easy on undirected graphs, and equivalent to FVS for directed graphs. In this paper, we deal with the vertex variant.

 DOI: 10.4086/toc.2016.v012a006

problem study *approximability* and *fixed-parameter tractability*. In fixed-parameter tractability, both undirected and directed FVS are shown to be in FPT [4, 5]. See recent results on a generalization of FVS [8, 6] and references therein. In this work, we focus on approximability.

FVS in undirected graphs has a 2-approximation algorithm [1, 3, 7], but the same problem is not well understood in directed graphs The best approximation algorithm [16, 11, 10] achieves an approximation factor of $O(\min(\log \tau^* \log\log \tau^*, \log n \log\log n))$, where τ^* is the optimal fractional solution in the natural LP relaxation.[2] The best hardness result follows from an easy approximation-preserving reduction from Vertex Cover by Dinur and Safra [9], which implies that it is NP-hard to approximate FVS within a factor of 1.36. Assuming the Unique Games Conjecture (UGC) of Khot [13], it is NP-hard (called *UG-hard*) to approximate FVS in directed graphs within any constant factor [12, 17].

The first UG-hardness of approximating FVS within a constant factor is the corollary of the fairly complicated result on the Maximum Acyclic Subgraph (MAS) by Guruswami et al. [12]. Svensson [17] gave a simpler proof tailored for FVS with a stronger statement in completeness—deleting $(1+\varepsilon)/k$ fraction of vertices ensures that there is no walk of length k. The main contribution of this work is a simpler proof of the same statement. Let $[i]_k$ be the integer in $\{1,\dots k\}$ such that $i \equiv [i]_k \bmod k$.

Theorem 1.1. *Fix an integer $k \geq 3$ and $\varepsilon \in (0, 1/(k+1))$. Given a directed graph $G = (V_G, E_G)$, it is UG-hard to distinguish the following cases.*

- *Completeness: The vertex set can be partitioned into sets $V_0, \dots, V_k$ such that*

$$|V_i| \geq \frac{(1-\varepsilon)}{k}|V_G|$$

  for all $i \in \{1, \dots, k\}$, and each edge not incident on V_0 goes from V_i to $V_{[i+1]_k}$ for some $i \in \{1, \dots, k\}$.

- *Soundness: Any subset of measure ε contains a k-cycle.*

Consequently, it is UG-hard to approximate FVS within a factor of k, for any constant k.

Our proof differs from Svensson's [17] in two aspects:

- The ingenious application of *It Ain't Over Till It's Over* Theorem is replaced by the standard application of the more general *Invariance principle* of Mossel [15].
- The reduction from Unique Games is simpler, introducing only one *long code* for each vertex of a Unique Games instance, while [17] used multiple long codes for each tuple of vertices of a certain length. Instead we rely on the stronger (but equivalent) UGC proposed by Khot and Regev [14].

The idea of using the Invariance principle to prove hardness of FVS is inspired by the elegant paper of Bansal and Khot [2] which showed structured hardness of k-Hypergraph Vertex Cover. Our main idea for this result is to use a more restricted distribution for the dictatorship test than the one used in [2] to ensure more structure in the completeness case. Our statement for the completeness case is stronger that of Theorem 1.1 of Svensson [17], but his technique also proves our statement. At the same time we also ensure that the distribution has certain properties so that the same soundness analysis can be applied.

[2]In unweighted cases, τ^* is always at most n. In weighted cases, we assume all weights are at least 1.

Notation. In a directed graph $G = (V_G, E_G)$, an edge (u, v) indicates a directed edge from u to v. In some cases G might be vertex-weighted or edge-weighted, and every weight will be normalized so that the sum is 1. Given a subset S of either V_G or E_G, define $\mu(S)$, also called the *measure* of S, to be the sum of the weights of the elements in S. Let $[k] := \{1, 2, \dots, k\}$. We often consider *hypercube* or *long code* $[k]^R$. We use superscripts $x^1, \dots, x^k \in [k]^R$ to denote k different points of the hypercube and subscripts $x_1, \dots, x_R$ to denote the value of each coordinate of one point $x \in [k]^R$.

Organization. In Section 2, we propose our *dictatorship test*. It is a family of instances of FVS where every small feedback vertex set must exhibit a certain structure, and the proposal and the analysis of the dictatorship test is our main technical contribution. Using the dictatorship test, Section 3 shows the full reduction from Unique Games to FVS, which is rather standard in the literature.

2 Dictatorship test

There is a simple gap-preserving reduction from FVS on vertex-weighted graphs to FVS on unweighted graphs—replace each vertex v by a set of new vertices $s(v)$ whose cardinality is proportional to the weight of v, and replace each edge (u, v) by $\{(u', v') : u' \in s(u), v' \in s(v)\}$. Our proof will have all the weights polynomially bounded, ensuring that this reduction runs in polynomial time. For the rest of the paper, we focus on vertex-weighted graphs.

We propose a simple dictatorship test for FVS, which is used to prove that it is UG-hard to approximate FVS within any constant factor. Given positive integers k, R, and $\varepsilon > 0$, our dictatorship test is a vertex-weighted graph $G = (V_G, E_G)$ where $V_G = ([k] \cup \{0\})^R$ and edges in E_G are carefully chosen to prove the following properties (informally stated).

- Completeness: For each $1 \le j \le R$, *depending only on the j-th coordinate*, V_G can be partitioned to $k + 1$ parts $V_0, \dots, V_k$ with the following two properties.
  - $\mu(V_0) = \varepsilon$, $\mu(V_1) = \cdots = \mu(V_k) = (1 - \varepsilon)/k$.
  - In the subgraph induced by $V_1 \cup \cdots \cup V_k$, each edge goes from V_i to $V_{[i+1]_k}$ for some $1 \le i \le k$.

  It is easy to see that $V_0 \cup V_i$ for any $1 \le i \le k$ gives a feedback vertex set with measure
  $$\varepsilon + \frac{1 - \varepsilon}{k}.$$
- Soundness: Any subset of measure at least ε that does not reveal any *influential coordinate* must contain a k-cycle.

Before defining G, we first define a k-uniform hypergraph $H = (V_H, E_H)$ with $V_H = V_G = (\{0\} \cup [k])^R$. The graph G is then simply obtained by replacing a hyperedge $(x^1, \dots, x^k)$ by k edges $(x^1, x^2), \dots, (x^k, x^1)$. The hypergraph H is vertex-weighted and edge-weighted. Both weights sum to 1 and induce probability distributions, where the weight of vertex x is the sum of the weight of the hyperedges containing x divided by k. The hyperedges of H are described by the following procedure to sample k vertices $(x^1, \dots, x^k)$ from $(\{0\} \cup [k])^R$, with the weight of each hyperedge equal to the probability that it is sampled in this procedure.

- For each coordinate $1 \leq j \leq R$, sample $(x^1)_j, \dots, (x^k)_j$ as follows, independently of the other coordinates.
  - Sample $a \in [k]$ uniformly at random.
  - Set $(x^1)_j = a, (x^2)_j = [a+1]_k, \dots, (x^k)_j = [a+k-1]_k$.
  - For each $(x^i)_j$, set $(x^i)_j = 0$ with probability ε independently.

This defines the hypergraph E_H. In the above distribution to sample $(x^1, \dots, x^k)$, the marginal on each x^i is the same:

$$\Pr[x^i = (a_1, \dots, a_R)] = \prod_{j=1}^{R} \mu(a_j) ,$$

where $\mu : [k] \cup \{0\} \to \mathbb{R}$ is defined by $\mu(0) = \varepsilon$ and $\mu(i) = (1-\varepsilon)/k$ for $i \in [k]$. Let the weight of $(x_1, \dots, x_R)$ be this quantity. The sum of the vertex weights is also 1.

With nonzero probability a randomly sampled hyperedge $(x^1, \dots, x^k)$ might have $x^i = x^j$ for some $i \neq j$. We call such hyperedges *defective* since they do not make H k-uniform. However, $x^i = x^j$ means $x^i = x^j = (0, 0, \dots, 0)$, so the probability that it happens is at most ε^{2R} and the sum of the weights of the defective hyperedges is at most $k^2 \varepsilon^{2R}$.

Finally, we define G. The vertex set $V_G = V_H$ with the same vertex weights, and for each non-defective hyperedge $(x^1, \dots, x^k) \in E_H$, we add k edges $(x^1, x^2), \dots, (x^k, x^1)$ to E_G. The analysis dealing with edge weights will be done in H, so we do not consider edge weights for the edges of G.

2.1 Analysis of dictatorship test

Completeness. Fix a coordinate $1 \leq j \leq R$. For all $0 \leq i \leq k$, let $V_i = \{(x_1, \dots, x_R) \in V_G : x_j = i\}$. By definition, $\mu(V_0) = \varepsilon, \mu(V_i) = (1-\varepsilon)/k$. The distribution on $(x^1, \dots, x^k)$ satisfies that for any $1 \leq i \leq k$, $(x^{[i+1]_k})_j = [(x^i)_j + 1]_k$ or at least one of $(x^i)_j, (x^{[i+1]_k})_j$ is 0. This proves that if we delete V_0 and the edges incident on it, all the remaining edges will go from V_i to $V_{[i+1]_k}$.

Soundness. We introduce some definitions and properties of correlated spaces and Fourier analysis of functions defined on (the products of) these spaces. See Mossel [15] for details.

Let $\Omega := [k] \cup \{0\}$ and $\mu : \Omega \to \mathbb{R}$ such that $\mu(0) = \varepsilon$ and $\mu(i) = (1-\varepsilon)/k$ as defined previously. Let (Ω^k, μ') be the probability space defined by the distribution of $(x^1)_j, \dots, (x^k)_j$ for some j from our hyperedge sampling. Note that the marginal distribution of each copy of Ω is μ. Given a probability space $(\Omega_1 \times \Omega_2, \nu)$, we define the correlation between Ω_1 and Ω_2 as

$$\rho(\Omega_1, \Omega_2; \nu) = \sup\left\{\mathsf{Cov}[f, g] : f \in \mathbb{R}^{\Omega_1}, g \in \mathbb{R}^{\Omega_2}, \mathsf{Var}[f] = \mathsf{Var}[g] = 1\right\} .$$

With more than two underlying spaces, the correlation of $(\Omega_1 \times \dots \times \Omega_k, \nu)$ is defined by

$$\rho(\Omega_1, \dots, \Omega_k; \nu) = \max_{1 \leq i \leq k} \rho\left(\prod_{j=1}^{i-1} \Omega_j \times \prod_{j=i+1}^{k} \Omega_j, \Omega_i; \nu\right) .$$

We use the following lemma to bound the correlation.

Lemma 2.1 (Lemma 2.9 of Mossel [15]). *Let $(\Omega_1 \times \Omega_2, \nu)$ be a probability space such that the probability of the smallest atom in $\Omega_1 \times \Omega_2$ is at least $\gamma > 0$. Define a bipartite graph $G = (\Omega_1 \cup \Omega_2, E)$ where $(a,b) \in \Omega_1 \times \Omega_2$ satisfies $(a,b) \in E$ if $\nu(a,b) > 0$. Then if G is connected then*

$$\rho(\Omega_1, \Omega_2; \nu) \leq 1 - \gamma^2/2\,.$$

In our distribution (Ω^k, μ'), note that $(0,0,\dots,0)$ has probability $\gamma := \varepsilon^k$, and this is indeed the smallest nonzero probability assuming $\varepsilon < 1/(k+1)$. Let $\Omega_1 = \Omega$, $\Omega_2 = \Omega^{k-1}$, and consider the bipartite graph defined above. For any $(x_1,\dots,x_k)$ with nonzero probability, the edge corresponding to $(x_1,\dots,x_k)$ is connected to the edge corresponding to $(0,0,\dots,0)$ since $(x_1,x_2,\dots,x_k)$, $(0,x_2,\dots,x_k)$, $(0,0,\dots,0)$ is the sequence of elements with nonzero probability where each consecutive elements differ in exactly one of Ω_1 or Ω_2 (i. e., consecutive edges share an endpoint in the bipartite graph). Therefore, we can apply the above lemma to see $\rho(\Omega_1, \Omega_2; \mu') \leq 1 - \gamma^2/2$. Since every copy of Ω is identical under μ', $\rho := \rho(\Omega,\dots,\Omega;\mu') \leq 1 - \gamma^2/2 < 1$. The similar argument also works for $\Omega_1 = \Omega^j$ and $\Omega_2 = \Omega^{k-j}$ for each $j \in \{1,\dots,k-1\}$, proving that

$$\rho(\Omega^j, \Omega^{k-j}, \mu') \leq 1 - \gamma^2/2\,, \quad \text{for each } j \in \{1,\dots,k-1\}.$$

Let $\chi_0,\dots,\chi_k \in \mathbb{R}^\Omega$ be orthonormal random variables satisfying that $\chi_0 \equiv 1$, $\mathbb{E}[\chi_i^2] = 1$ for all i, and $\mathbb{E}[\chi_i \chi_j] = 0$ for all $i \neq j$. Given $f : \Omega^R \to [0,1]$ as a random variable in the probability space $(\Omega^R, \mu^{\otimes R})$, its multilinear decomposition is

$$f(x_1,\dots,x_R) = \sum_{\alpha \in \Omega^R} \hat{f}(\alpha) \prod_{j=1}^{R} \chi_{\alpha(j)}(x_j)\,.$$

Let $\mathsf{Supp}(\alpha)$ be the number of nonzero coordinates of α. The d-degree influence of the j-th coordinate of f is defined by

$$\mathsf{Inf}_j^{\leq d}(f) = \sum_{\alpha \in \Omega^R : \alpha_j \neq 0, \mathsf{Supp}(\alpha) \leq d} \hat{f}(\alpha)^2.$$

It is well known that $\sum_{j=1}^R \mathsf{Inf}_j^{\leq d}(f) \leq d$ for $[0,1]$-valued f and does not depend on the choice of $\chi_0,\dots,\chi_k$.

We establish the soundness property using the Invariance principle stated below.

Theorem 2.2 (Theorem 6.3 of Mossel [15]). *Let $(\prod_{i=1}^k \Omega_i, \nu)$ be a probability space such that the minimum probability of any atom is at least $\gamma > 0$. Assume furthermore that there exists $\rho < 1$ such that*

$$\rho(\Omega_1,\dots,\Omega_k,\nu) \leq \rho\,,$$
$$\rho\Big(\prod_{l=1}^{i} \Omega_l, \prod_{l=i+1}^{k} \Omega_l, \nu\Big) \leq \rho, \quad \textit{for all } i \in \{1,\dots,k-1\}\,.$$

Then for all $\beta > 0$, there exist $\delta > 0, \tau > 0$, and an integer d such that the following holds. Fix a natural number R and consider the space $(\prod_{i=1}^k \Omega_i^R, \nu^{\otimes R})$. If k functions $\{f_i : \Omega_i^R \mapsto [0,1]\}_{1 \leq i \leq k}$ satisfy

$$\mathbb{E}[f_i] \geq \beta\,, \qquad i \in [k]\,,$$
$$\mathsf{Inf}_j^{\leq d}(f_i) \leq \tau\,, \qquad \forall i \in [k], j \in [R]\,,$$

then

$$\mathbb{E}\left[\prod_{i=1}^{k} f_j\right] \geq \delta.$$

The original statement of Mossel [15] is more general than the above statement and lower bounds $\mathbb{E}[\prod_{i=1}^{k} f_i]$ by some quantity $\Gamma := \Gamma(\rho,\beta)$ minus some additive error. Given $\rho < 1$ and $\beta > 0$, our statement is obtained by observing that $\Gamma(\rho,\beta) > 0$ and setting the additive error to be $\Gamma(\rho,\beta)/2$ so that $\delta := \Gamma(\rho,\beta)/2$ becomes a lower bound of $\mathbb{E}[\prod_{j=1}^{k} f_j]$. We refer the reader to [15] for the original statement.

Let A be the subset of V_G of measure at least β, and f be its indicator function. Apply Theorem 2.2 with ρ, β, and $\nu \leftarrow \mu'$ to have δ, τ and d. If $\mathsf{Inf}_j(f) \leq \tau$ for all $j \in [R]$ (i. e., A does not reveal any influential coordinate), as long as δ is greater than the sum of the weights of the defective hyperedges, which is at most $k^2\varepsilon^{2R}$ (which can be ensured by taking large R for fixed k and ε), A contains a non-defective hyperedge $(x^1,\ldots,x^k)$ of H and the corresponding k-cycle of G. By taking $\beta \leftarrow \varepsilon$, we can conclude that any subset of measure at least ε that does not reveal any influential coordinate must contain a k-cycle, establishing the desired soundness property.

3 Reduction from the Unique Games

We introduce the Unique Games Conjecture and its equivalent variant.

Definition 3.1. An instance

$$\mathcal{L}\left(B(V_B \cup W_B, E_B), [R], \{\pi(v,w)\}_{(v,w)\in E_B}\right)$$

of Unique Games consists of a biregular bipartite graph $B(V_B \cup W_B, E_B)$ and a set $[R]$ of labels. For each edge $(v,w) \in E_B$ there is a constraint specified by a permutation $\pi(v,w) : [R] \to [R]$. The goal is to find a labeling $\ell : V_B \cup W_B \to [R]$ of the vertices such that as many edges as possible are satisfied, where an edge $e = (v,w)$ is said to be satisfied if $\ell(v) = \pi(v,w)(\ell(w))$.

Definition 3.2. Given a Unique Games instance

$$\mathcal{L}\left(B(V_B \cup W_B, E_B), [R], \{\pi(v,w)\}_{(v,w)\in E_B}\right),$$

let $\mathsf{Opt}(\mathcal{L})$ denote the maximum fraction of simultaneously-satisfied edges of $\mathcal{L}$ by any labeling, i.e.,

$$\mathsf{Opt}(\mathcal{L}) := \frac{1}{|E|} \max_{\ell : V_B \cup W_B \to [R]} |\{e \in E : \ell \text{ satisfies } e\}|.$$

Conjecture 3.3 (The Unique Games Conjecture [13])**.** *For any constants $\eta > 0$, there is $R = R(\eta)$ such that, for a Unique Games instance $\mathcal{L}$ with label set $[R]$, it is* NP*-hard to distinguish between the following cases.*

- $\mathsf{opt}(\mathcal{L}) \geq 1 - \eta$.

- $\mathsf{opt}(\mathcal{L}) \le \eta$.

To show the optimal hardness result for Vertex Cover, Khot and Regev [14] introduced the following seemingly stronger conjecture, and proved that it is in fact equivalent to the original Unique Games Conjecture.

Conjecture 3.4 (Khot and Regev [14]). *For any constants $\eta > 0$, there is $R = R(\eta)$ such that, for a Unique Games instance $\mathcal{L}$ with label set $[R]$, it is NP-hard to distinguish between the following cases.*

- *There is a set $W' \subseteq W_B$ such that $|W'| \ge (1-\eta)|W_B|$ and a labeling $\ell : V_B \cup W_B \to [R]$ that satisfies every edge (v,w) for $v \in V_B$ and $w \in W'$.*
- $\mathsf{opt}(\mathcal{L}) \le \eta$.

We describe the reduction from Unique Games. It is parametrized by an integer k and $\varepsilon \in (0, 1/(k+1))$ as in the statement of Theorem 1.1 and another parameter R that will be chosen later. Note that k and ε determine the correlated space (Ω^k, μ') as in the previous section.

Given an instance $\mathcal{L}$ of Unique Games, we assign to each vertex $w \in W_B$ the hypercube Ω^R_w. Formally, $V_G = V_H := W_B \times \Omega^R$. The weight of each vertex (w,x) is the weight of x in Ω^R divided by $|W_B|$, so that the sum of the weights is again 1.

For a permutation $\sigma : [R] \to [R]$, let $x \circ \sigma := (x_{\sigma(1)}, \dots, x_{\sigma(R)})$. The weighted hyperedges of H are again defined by the following procedure to sample k vertices $(w^1, x^1), \dots, (w^k, x^k)$.

- Sample $v \in V_B$ uniformly at random.
- Sample k vertices $w^1, \dots, w^k \in W_B$ i. i. d. from neighbors of v.
- Sample $x^1, \dots, x^k \in \Omega^R$ from the dictatorship distribution.
- Return the hyperedge $((w^1, x^1 \circ \pi(v, w^1)), \dots, (w^k, x^k \circ \pi(v, w^k)))$.

For each non-defective hyperedge $((w^1,x^1), \dots, (w^k,x^k))$, we add k edges

$$((w^1,x^1),(w^2,x^2)), \dots, ((w^k,x^k),(w^1,x^1))$$

to G.

Completeness. Suppose there exists a labeling ℓ and a subset $W' \subseteq W_B$ with $|W'| \ge (1-\eta)|W_B|$ such that ℓ satisfies every edge incident on W'. For $1 \le i \le k$, let

$$V_i := \bigcup_{w \in W'} \{(w,x) : x_{\ell(w)} = i\}$$

and $V_0 := V_G \setminus (\bigcup_{i=1}^k V_i)$. Note that for $i \in [k]$, $\mu(V_i) \ge (1-\eta)(1-\varepsilon)/k$. Let G' be the induced subgraph on $V_G \setminus V_0$. For any edge $((w^1,x^1),(w^2,x^2)) \in E_{G'}$, we know $w^1, w^2 \in W'$ and they share a neighbor $v \in V_B$. By the property of our dictatorship test, for each $1 \le j \le R$,

$$a := (x^1)_{\pi(v,w^1)^{-1}(j)} \quad \text{and} \quad b := (x^2)_{\pi(v,w^2)^{-1}(j)}$$

satisfy that at least one of them is zero or $b = [a+1]_k$. Therefore, if $(w^1,x^1),(w^2,x^2) \notin V_0$, which implies

$$(x^1)_{\pi(v,w^1)^{-1}(\ell(v))} = (x^1)_{\ell(w^1)}, \quad (x^2)_{\pi(v,w^2)^{-1}(\ell(v))} = (x^2)_{\ell(w^2)}$$

are nonzero, we can conlude that $(w^1,x^1) \in V_i$ and $(w^2,x^2) \in V_{[i+1]_k}$ for some $1 \le i \le k$.

Soundness. The soundness anlaysis is standard and closely follows Bansal and Khot [2]. Suppose $A \subseteq V_H$ of measure at least β such that it is independent (i.e., does not contain any non-defective hyperedge). We will show that the instance $\mathcal{L}$ of Unique Games admits a good labeling. Its contrapositive shows that if $\mathcal{L}$ does not admit a good labeling, any subset of measure at least β contains a non-defect hyperedge and the corresponding k-cycle, proving Theorem 1.1.

Let $A_w = \Omega_w^R \cap A$ be the vertices of A that lie in Ω_w^R for $w \in W_B$. Let $f_w : \Omega^R \to \{0,1\}$ be the indicator function of A_w. Define $f_v : \Omega^R \to [0,1]$ for each $v \in V_B$ to be

$$f_v(x) = \mathbb{E}_{w \in N(v)}[f_w(x \circ \pi(v,w))]$$

where $N(v)$ is the set of neighbors of v. Since B is biregular, $\mathbb{E}_{v,x}[f_v(x)] \geq \beta$. By an averaging argument, at least $\beta/2$ fraction of vertices in V_B satisfy $\mathbb{E}_x[f_v(x)] \geq \beta/2$. Call such vertices *good*.

Since A is an independent set, for any $v \in V$ and its k neighbors $w^1, \ldots, w^k$, we have

$$\mathbb{E}_{x^1,\ldots,x^k}\left[\prod_{i=1}^{k} f_{w^i}(x^i \circ \pi(v,w^i))\right] \leq k^2 \varepsilon^{2R}.$$

Averaging over all k-tuples $w^1, \ldots, w^k$ of neighbors of v, we have

$$\mathbb{E}_{x^1,\ldots,x^k}\left[\prod_{i=1}^{k} f_v(x^i)\right] = \mathbb{E}_{x^1,\ldots,x^k}\mathbb{E}_{w^1,\ldots,w^k \in N(v)}\left[\prod_{i=1}^{k} f_{w^i}(x^i \circ \pi(v,w^i))\right] \leq k^2 \varepsilon^{2R}.$$

Applying Theorem 2.2 (take R large enough to make sure that $k^2\varepsilon^{2R} \ll \delta$), there exist τ and d such that f_v has a coordinate j with $\mathsf{Inf}_j^{\leq d}(f_v) \geq \tau$. Set $\ell(v) = j$. Since

$$\begin{aligned}\mathsf{Inf}_j^{\leq d}(f_v) &= \sum_{\alpha_j \neq 0, |\alpha| \leq d} \hat{f}_v(\alpha)^2 = \sum_{\alpha_j \neq 0, |\alpha| \leq d} \left(\mathbb{E}_w[\hat{f}_w(\pi(v,w)^{-1}(\alpha))]^2\right) \\ &\leq \sum_{\alpha_j \neq 0, |\alpha| \leq d} \mathbb{E}_w\left[\hat{f}_w(\pi(v,w)^{-1}(\alpha))^2\right] = \mathbb{E}_w\left[\mathsf{Inf}^{\leq d}_{\pi(v,w)^{-1}(j)}(f_w)\right],\end{aligned}$$

at least $\tau/2$ fraction of v's neighbors satisfy $\mathsf{Inf}^{\leq d}_{\pi(v,w)^{-1}(j)}(f_w) \geq \tau/2$. There are at most $2d/\tau$ coordinates with degree-d influence at most $\tau/2$, and $\ell(w)$ is chosen uniformly among those coordinates (if there is none, set it arbitrarily). The above probabilistic strategy satisfies at least $(\beta/2)(\tau/2)(\tau/2d)$ fraction of all edges. By taking large R, η can be made less than this quantity, implying that if a Unique Games instance has value at most η, then the resulting H cannot have an independent set of measure at least β, which is equivalent to saying that every subset of V_G of measure at least β contains a k-cycle. Taking $\beta \leftarrow \varepsilon$ proves Theorem 1.1.

References

[1] VINEET BAFNA, PIOTR BERMAN, AND TOSHIHIRO FUJITO: Constant ratio approximations of the weighted feedback vertex set problem for undirected graphs. In *Proc. 6th Internat. Symp. Alg. Comput. (ISAAC'95)*, pp. 142–151. Springer, 1995. [doi:10.1007/BFb0015417] 2

[2] NIKHIL BANSAL AND SUBHASH KHOT: Inapproximability of hypergraph vertex cover and applications to scheduling problems. In *Proc. 37th Internat. Colloq. on Automata, Languages and Programming (ICALP'10)*, pp. 250–261. Springer, 2010. [doi:10.1007/978-3-642-14165-2_22] 2, 8

[3] ANN BECKER AND DAN GEIGER: Optimization of Pearl's method of conditioning and greedy-like approximation algorithms for the vertex feedback set problem. *Artificial Intelligence*, 83(1):167–188, 1996. [doi:10.1016/0004-3702(95)00004-6] 2

[4] HANS L. BODLAENDER: On disjoint cycles. *Internat. J. Found. Comput. Sci.*, 5(1):59–68, 1994. Preliminary version in WG'91. [doi:10.1142/S0129054194000049] 2

[5] JIANER CHEN, YANG LIU, SONGJIAN LU, BARRY O'SULLIVAN, AND IGOR RAZGON: A fixed-parameter algorithm for the directed feedback vertex set problem. *J. ACM*, 55(5):21:1–21:19, 2008. Preliminary version in STOC'08. [doi:10.1145/1411509.1411511] 2

[6] RAJESH CHITNIS, MAREK CYGAN, MOHAMMADTAGHI HAJIAGHAYI, AND DÁNIEL MARX: Directed subset feedback vertex set is fixed-parameter tractable. *ACM Trans. Algorithms*, 11(4):28:1–28:28, 2015. Preliminary version in ICALP'12. [doi:10.1145/2700209, arXiv:1205.1271] 2

[7] FABIÁN A. CHUDAK, MICHEL X. GOEMANS, DORIT S. HOCHBAUM, AND DAVID P. WILLIAMSON: A primal-dual interpretation of two 2-approximation algorithms for the feedback vertex set problem in undirected graphs. *Oper. Research Lett.*, 22(4-5):111–118, 1998. [doi:10.1016/S0167-6377(98)00021-2] 2

[8] MAREK CYGAN, MARCIN PILIPCZUK, MICHAŁ PILIPCZUK, AND JAKUB ONUFRY WOJTASZCZYK: Subset feedback vertex set is fixed-parameter tractable. *SIAM J. Discrete Math.*, 27(1):290–309, 2013. Preliminary version in ICALP'11. [doi:10.1137/110843071, arXiv:1004.2972] 2

[9] IRIT DINUR AND SHMUEL SAFRA: On the hardness of approximating minimum vertex cover. *Ann. of Math.*, 162(1):439–485, 2005. Preliminary version in STOC'02. [doi:10.4007/annals.2005.162.439] 2

[10] GUY EVEN, JOSEPH SEFFI NAOR, SATISH RAO, AND BARUCH SCHIEBER: Divide-and-conquer approximation algorithms via spreading metrics. *J. ACM*, 47(4):585–616, 2000. Preliminary version in FOCS'95. [doi:10.1145/347476.347478] 2

[11] GUY EVEN, JOSEPH SEFFI NAOR, BARUCH SCHIEBER, AND MADHU SUDAN: Approximating minimum feedback sets and multicuts in directed graphs. *Algorithmica*, 20(2):151–174, 1998. Preliminary version in IPCO'95. [doi:10.1007/PL00009191] 2

[12] VENKATESAN GURUSWAMI, JOHAN HÅSTAD, RAJSEKAR MANOKARAN, PRASAD RAGHAVENDRA, AND MOSES CHARIKAR: Beating the random ordering is hard: Every ordering CSP is approximation resistant. *SIAM J. Comput.*, 40(3):878–914, 2011. Preliminary versions in FOCS'08 and ECCC. [doi:10.1137/090756144] 2

[13] SUBHASH KHOT: On the power of unique 2-prover 1-round games. In *Proc. 34th STOC*, pp. 767–775. ACM Press, 2002. [doi:10.1145/509907.510017] 2, 6

[14] SUBHASH KHOT AND ODED REGEV: Vertex cover might be hard to approximate to within $2-\varepsilon$. *J. Comput. System Sci.*, 74(3):335–349, 2008. Preliminary version in CCC'03. [doi:10.1016/j.jcss.2007.06.019] 2, 7

[15] ELCHANAN MOSSEL: Gaussian bounds for noise correlation of functions. *Geom. Funct. Anal.*, 19(6):1713–1756, 2010. Preliminary version in FOCS'08. [doi:10.1007/s00039-010-0047-x, arXiv:math/0703683] 2, 4, 5, 6

[16] PAUL D. SEYMOUR: Packing directed circuits fractionally. *Combinatorica*, 15(2):281–288, 1995. [doi:10.1007/BF01200760] 2

[17] OLA SVENSSON: Hardness of vertex deletion and project scheduling. *Theory of Computing*, 9(24):759–781, 2013. Preliminary version in APPROX'12. [doi:10.4086/toc.2013.v009a024, arXiv:1206.3408] 2

AUTHORS

Venkatesan Guruswami
Professor
Carnegie Mellon University, Pittsburgh, PA
guruswami@cmu.edu
`http://www.cs.cmu.edu/~venkatg`

Euiwoong Lee
Graduate Student
Carnegie Mellon University, Pittsburgh, PA
euiwoonl@cs.cmu.edu
`http://www.cs.cmu.edu/~euiwoonl`

ABOUT THE AUTHORS

VENKATESAN GURUSWAMI, or Venkat, as his colleagues call him, received his Bachelor's degree from the Indian Institute of Technology at Madras in 1997. Venkat received his Ph. D. from the Massachusetts Institute of Technology in 2001. He was fortunate to be advised by Madhu Sudan, with whom he enjoyed collaboration on a number of topics, most notably on list decoding.

Since 2008, Venkat has been on the faculty of the Computer Science Department at Carnegie Mellon University. Earlier, he was a faculty member at the University of Washington. Venkat was a Miller Research Fellow at UC Berkeley during 2001-02, a member of the School of Mathematics, Institute for Advanced Study during 2007-08, and a visiting researcher at Microsoft Research New England during January-June 2014.

Venkat is interested in several topics in theoretical computer science, including algorithmic and algebraic coding theory, approximability of fundamental optimization problems, pseudorandomness, PCPs, and computational complexity theory.

Venkat currently serves on the editorial boards of the Journal of the ACM, the SIAM Journal on Computing, and the journal Research in the Mathematical Sciences, and was previously on the editorial boards of the IEEE Transactions on Information Theory, and the ACM Transactions on Computation Theory. He was recently Program Committee chair for the 2015 Foundations of Computer Science conference and the 2012 Computational Complexity Conference. Venkat is a recipient of the Presburger award, Packard Fellowship, Sloan Fellowship, NSF CAREER award, ACM's Doctoral Dissertation Award, and the IEEE Information Theory Society Paper Award.

In his (sadly limited) spare time, Venkat enjoys traveling, ethnic vegetarian food, racquet sports, and listening to Carnatic music.

EUIWOONG LEE is a Ph. D. student in computer science at Carnegie Mellon University, advised by Venkat Guruswami. He is interested in various aspects of approximation algorithms, including hardness of approximation and the power of LP / SDP hierarchies. After graduating from high school in Seoul, South Korea, he moved to the U. S. and received a B. S. and an M. S. degrees from Caltech. He enjoys playing and watching basketball.