THEORY OF COMPUTING, Volume 11 (7), 2015, pp. 221–235
www.theoryofcomputing.org

SPECIAL ISSUE: APPROX-RANDOM 2012

The Projection Games Conjecture and the NP-Hardness of $\ln n$-Approximating SET-COVER

Dana Moshkovitz*

Received October 21, 2012; Revised July 22, 2014; Published June 9, 2015

Abstract: We establish a tight NP-hardness result for approximating the SET-COVER problem based on a strong PCP theorem. Our work implies that it is NP-hard to approximate SET-COVER on instances of size N to within $(1-\alpha)\ln N$ for arbitrarily small $\alpha > 0$. Our reduction establishes a tight trade-off between the approximation accuracy α and the running time $\exp(N^{\Omega(\alpha)})$ assuming SAT requires exponential time.

The reduction is obtained by modifying Feige's reduction. The latter provides a lower bound of $\exp(N^{\Omega(\alpha/\log\log N)})$ on the time required for $(1-\alpha)\ln N$-approximating SET-COVER assuming SAT requires exponential time. The modification uses a combinatorial construction of a bipartite graph in which any coloring of the first side that does not use a color for more than a small fraction of the vertices, makes most vertices on the other side have all their neighbors colored in different colors.

In the conference version of this paper, the SET-COVER result was conditioned on a conjecture we call "The Projection Games Conjecture" (PGC), a strengthening of the Sliding

A preliminary version of this paper appeared in the Proceedings of The 15th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2012) [30].

*This material is based upon work supported by the National Science Foundation under Grant Number 1218547.

ACM Classification: F.2.2, F.1.3, G.1.6, F.2.3

AMS Classification: 68Q17, 68W25, 68Q25

Key words and phrases: Set-Cover, PCP, Sliding Scale Conjecture, Projection Games Conjecture

DOI: 10.4086/toc.2015.v011a007

Scale Conjecture of Bellare, Goldwasser, Lund and Russell to projection games (LABEL-COVER). More precisely, the prerequisite was a quantitative version of this conjecture that was slightly beyond what was known at the time of the paper's writing. Shortly afterward, Dinur and Steurer, based on a result by the author and Raz, proved the quantitative version of the conjecture sufficient for the SET-COVER result. More broadly, in this paper we discuss the Projection Games Conjecture and its applications to hardness of approximation, e. g., to polynomial hardness factors for the CLOSEST-VECTOR problem and to studying the behavior of CSPs around their approximability threshold.

1 SET-COVER

In SET-COVER, given a collection of subsets of a base set such that the sets cover all of the base set, the goal is to find as few of the sets as possible that cover the entire base set.

Definition 1.1 (SET-COVER)**.** The input to SET-COVER consists of a base set U, $|U| = n$ and subsets $S_1, \ldots, S_m \subseteq U$, $\bigcup_{j=1}^{m} S_j = U$. The goal is to find as few sets $S_{i_1}, \ldots, S_{i_k}$ as possible that cover U, i. e., $\bigcup_{j=1}^{k} S_{i_j} = U$.

SET-COVER is a classical NP-hard optimization problem. It is equivalent to the HITTING-SET, HYPERGRAPH-VERTEX-COVER and DOMINATING-SET problems, and is a special case of many other problems, e. g., GROUP-STEINER-TREE and GROUP-TRAVELING-SALESMAN-PROBLEM.

The greedy algorithm was shown to give a $(\ln n + 1)$-approximation for SET-COVER [19, 25, 42]. Slavík analyzed the low-order terms of the greedy algorithm, and showed that it in fact obtains an approximation to within $\ln n - \ln\ln n + O(1)$ [40]. SET-COVER also has a linear programming based algorithm that gives approximation to within the same factor [41].

Lund and Yannakakis proved that SET-COVER cannot be approximated in polynomial time to within any factor better than $(\log_2 n)/4$, assuming $\mathsf{NP} \not\subseteq \mathsf{DTIME}(n^{\text{poly}\log n})$ [26]. By adapting their construction, Feige changed the leading constant to the right constant, and showed that SET-COVER cannot be approximated in polynomial time to within $(1-\alpha)\ln n$ for any $\alpha > 0$, assuming $\mathsf{NP} \not\subseteq \mathsf{DTIME}(n^{O(\lg\lg n)})$ [14]. (The improvement in the assumption is due to the parallel repetition theorem [36], proved in the time between the two results.) Under $\mathsf{P} \neq \mathsf{NP}$, the best hardness factor known prior to this work was about $0.2\ln n$ [2], based on the PCP of [37, 6].

Parallel repetition is used by Feige not only to ensure very low error, $1/(\log n)^{O(1)}$, for the PCP, but also for its unique structure. It was assumed by some that the blow-up incurred by parallel repetition was inherent to SET-COVER. We show that this is not the case. The following theorem follows from the reduction presented in this paper together with the parallel repetition of Dinur and Steurer [13].

Theorem 1.2. *For every* $0 < \alpha < 1$*, (exact)* SAT *on inputs of size n can be reduced in polynomial time to approximating* SET-COVER *to within* $(1-\alpha)\ln N$ *on inputs of size* $N = n^{O(1/\alpha)}$.

The theorem shows that approximating SET-COVER on inputs of size N better than $(1-\alpha)\ln N$ is NP-hard. Interestingly, the $N = n^{O(1/\alpha)}$ blow-up of the reduction is optimal (up to the constant in the $O(\cdot)$), assuming that SAT requires exponential time, $2^{\Omega(n)}$ ("The Exponential Time Hypothesis" [18]). This follows from a slightly subexponential $2^{O(N^{\alpha}+\text{poly}\log N)}$-time approximation algorithm for $(1-\alpha)\ln N$ approximating SET-COVER [10].

2 Projection games and the Projection Games Conjecture

In the conference version of this paper [30], Theorem 1.2 was conditioned on a conjecture we call "The Projection Games Conjecture" (PGC), or, more precisely, on a quantitative version of this conjecture that was slightly beyond what was known at the time of the paper's writing. Shortly afterward, Dinur and Steurer [13], based on a result by the author and Raz [31], proved the quantitative version of the conjecture sufficient for Theorem 1.2. In this section we discuss the Projection Games Conjecture.

Most of the NP-hardness of approximation results known today (e. g., all of the results in Håstad's paper [16]) are based on a PCP Theorem for LABEL-COVER. The input to LABEL-COVER consists of (i) a bipartite graph $G = (A, B, E)$; (ii) finite alphabets Σ_A, Σ_B; (iii) constraints (also called *projections*) $\pi_e : \Sigma_A \to \Sigma_B$ for every edge $e \in E$. The goal is to find assignments to the vertices, $\varphi_A : A \to \Sigma_A$, $\varphi_B : B \to \Sigma_B$, that *satisfy* as many of the edges as possible. We say that an edge $e = (a, b) \in E$ is satisfied if the projection constraint holds, i. e., $\pi_e(\varphi_A(a)) = \varphi_B(b)$. We denote the size of the label cover by $n = |A| + |B| + |E|$. The size of the alphabet is $\max\{|\Sigma_A|, |\Sigma_B|\}$.

A PCP Theorem for LABEL-COVER with soundness error ε and alphabet size k (where ε and k may depend on n) states the following [5, 4, 36]:

> *It is* NP-*hard, given an input of size n for* LABEL-COVER *with alphabets of size k, to distinguish between the case where all edges can be satisfied and the case where at most ε fraction of the edges can be satisfied.*

We can refine this statement by saying that there is a reduction from (exact) SAT to LABEL-COVER, which maps instances of SAT of size n to instances of LABEL-COVER of size $N = n^{1+o(1)}\operatorname{poly}(1/\varepsilon)$. Such PCPs are referred to as "almost-linear size PCP" because of the exponent of n, although for small ε the blow-up may be super-linear.

Ran Raz and the author proved the following result.

Theorem 2.1 ([32])**.** *There exists $c > 0$, such that for every $\varepsilon \geq 1/N^c$,* SAT *on input of size n can be reduced to* LABEL-COVER *of size N for $N = n^{1+o(1)}\operatorname{poly}(1/\varepsilon)$. The* LABEL-COVER *is over an alphabet of size exponential in $1/\varepsilon$, and has soundness error ε. The reduction can be computed in linear time in the size and the alphabet size of the* LABEL-COVER *instance. The* LABEL-COVER *is on a bi-regular graph whose degrees are* $\operatorname{poly}(1/\varepsilon)$.

One cannot hope for a soundness error that is lower than $1/N$. Hence, the dependence of ε on N is as low as possible up to the value of the constant c. On the other hand, the alphabet size in Theorem 2.1 is not known to be tight. It can be shown that the alphabet size must be at least $1/\varepsilon$ where ε is the soundness error (assuming $\mathsf{P} \neq \mathsf{NP}$). Moreover, certain PCP constructions—while deficient in other parameters—have alphabets of size $\operatorname{poly}(1/\varepsilon)$, see, e. g., [36]. This motivates the conjecture that an alphabet size of $\operatorname{poly}(1/\varepsilon)$ could be achieved in Theorem 2.1 as well.

Conjecture 2.2 (Projection Games Conjecture,[1] PGC)**.** *There exists $c > 0$, such that for every $\varepsilon \geq 1/N^c$,* SAT *on inputs of size n can be efficiently reduced to* LABEL-COVER *of size $N = n^{1+o(1)}\operatorname{poly}(1/\varepsilon)$ over an alphabet of size* $\operatorname{poly}(1/\varepsilon)$ *that has soundness error ε.*

[1] A slightly weaker version of the Projection Games Conjecture is one in which the size of the label cover is polynomial, $N = \operatorname{poly}(n, 1/\varepsilon)$, rather than almost-linear.

In almost all applications, one wishes the size and the alphabet size of the LABEL-COVER to be at most polynomial in n. This happens in Theorem 2.1 only when $\varepsilon \geq 1/(\log N)^b$ for a constant $b > 0$. The PGC, on the other hand, would give polynomial size and polynomial alphabet size for any $\varepsilon \geq 1/N^c$.

The PGC is the strengthening of the Sliding Scale Conjecture of Bellare, Goldwasser, Lund and Russell [7] obtained by restricting it to LABEL-COVER (of almost-linear size). "Sliding scale" refers to the idea that the error can be decreased as we increase the alphabet size. Bellare et al. conjectured that polynomially small error could be achieved simultaneously with polynomial alphabet, even for two queries. They did not formulate their conjecture for LABEL-COVER, which was introduced by Arora et al. [3] around the same time. Today, focusing on the PGC became natural in the PCP community.

Approximation algorithms for LABEL-COVER were designed [34, 9, 27], and the conjecture is consistent with the state of the art algorithm, giving $1/\varepsilon = O(\sqrt[4]{Nk})$ [27]. For PCPs with more than two queries (corresponding to games on hypergraphs, where the edges carry general predicates rather than projections), soundness error approaching polynomial, $\varepsilon = 2^{-(\log N)^{1-\alpha}}$ for every $\alpha > 0$, is known [11]. Alas, these PCPs are not for label covers, and the number of queries depends on ε.

Dinur and Steurer show how to achieve soundness error that is poly-logarithmic in N (for *any* poly-logarithm) simultaneously with polynomial-sized alphabet, at the cost of increasing the size. This suffices for the reduction to SET-COVER in Theorem 1.2 to go through. The idea is to apply parallel repetition on Theorem 2.1, and Dinur and Steurer were the first to successfully analyze parallel repetition for the relevant parameters.

Theorem 2.3 ([13]). *There exists $c > 0$, such that for every $\varepsilon \geq 1/N^c$ and every $k \geq 1$,* SAT *on input of size n can be reduced to* LABEL-COVER *on a bi-regular graph whose size is N^k for $N = n^{1+o(1)}\operatorname{poly}(1/\varepsilon)$. The projection game is over an alphabet of size exponential in $1/\varepsilon$ and k, and has soundness error $(2\varepsilon)^{k/2}$. The reduction can be computed in linear time in the size and the alphabet size of the* LABEL-COVER. *The* LABEL-COVER *is on a bi-regular graph whose degrees are* $\operatorname{poly}(1/\varepsilon^k)$.

The Projection Games Conjecture has a similar flavor to Khot's Unique Games Conjecture (UGC) [21]; both assert that low soundness error[2] for a special kind of 2-prover games can be obtained for sufficiently large alphabets. Unique games are the special case of LABEL-COVER in which the projections π_e are one-to-one. Unique games are easier than general projection games. In particular, while there are constructions of projection games with low soundness error for SAT, we do not know of any constructions of unique games with almost-perfect completeness[3] and bounded soundness error. The two conjectures, UGC and PGC, seem unrelated: neither would imply the other.

The following variant of the PGC is useful for hardness of approximation.

Definition 2.4 (Linear projection game). A linear LABEL-COVER is LABEL-COVER where the alphabets are of the form $\Sigma_A = \mathbb{F}^a$, $\Sigma_B = \mathbb{F}^b$, where $\mathbb{F}$ is a finite field, and $a \geq b$ are natural numbers. The projections in the game are affine transformations $\mathbb{F}^a \to \mathbb{F}^b$.

[2]The unique games conjecture only asks for arbitrarily small constant soundness error ε, while the PGC asks for polynomially small error.

[3]For unique games, if all the edges can be satisfied simultaneously, then one can find a satisfying assignment in polynomial time. Hence, we consider the case where *almost* all edges can be satisfied simultaneously ("almost perfect completeness").

Conjecture 2.5 (Linear PGC)**.** *There exists $c > 0$, such that for every $\varepsilon \geq 1/n^c$,* SAT *on inputs of size n can be efficiently reduced to a linear* LABEL-COVER *of size $N = n^{1+o(1)}\,\mathrm{poly}(1/\varepsilon)$ and alphabet size* $\mathrm{poly}(1/\varepsilon)$. *Satisfiable instances of* SAT *are mapped to* LABEL-COVER *where $1-\varepsilon$ fraction of the edges can be satisfied, while unsatisfiable instances of* SAT *are mapped to* LABEL-COVER *where at most ε fraction of the edges can be satisfied.*

Note that for linear LABEL-COVER, one can efficiently distinguish the case where all edges can be satisfied from the case where not all edges can be satisfied—by Gaussian elimination. Therefore, it was necessary to modify the statement of Conjecture 2.2.

In Section 5 we discuss applications of the PGC and the linear PGC to proving polynomial hardness factors for the CLOSEST-VECTOR-PROBLEM and to studying the behavior of MAX-3LIN and other CSPs around their approximability thresholds.

3 Preliminaries

For a set S and a natural number ℓ we denote by $\binom{S}{\ell}$ the family of all sets of ℓ elements from S.

We assume without loss of generality that the LABEL-COVER instance in Conjecture 2.2 is bi-regular, i. e., all the vertices from A have the same degree, which we call the A-degree, and all the vertices from B have the same degree, which we call the B-degree. We note that any LABEL-COVER instance can be converted to bi-regular using a technique developed in [32] ("right degree reduction—switching sides—right degree reduction"), and the cost in the soundness error and graph size does not change the parameters as stated in Conjecture 2.2.

4 SET-COVER hardness

4.1 The new component

Feige uses the structure obtained from parallel repetition to achieve a LABEL-COVER in which the soundness guarantee is that very few vertices from B have any two of their neighbors agree on a value for them.

Definition 4.1 (Total disagreement)**.** Let $(G = (A,B,E),\Sigma_A,\Sigma_B,\Phi)$ be a LABEL-COVER instance. Let $\varphi_A : A \to \Sigma_A$ be an assignment to the A-vertices. We say that the A-vertices *totally disagree* on a vertex $b \in B$ if there are no two neighbors $a_1, a_2 \in A$ of b, for which

$$\pi_{e_1}(\varphi_A(a_1)) = \pi_{e_2}(\varphi_A(a_2)),$$

where $e_1 = (a_1,b), e_2 = (a_2,b) \in E$.

Definition 4.2 (Agreement soundness)**.** Let $\mathcal{G} = (G = (A,B,E),\Sigma_A,\Sigma_B,\Phi)$ be a LABEL-COVER for deciding whether a Boolean formula ϕ is satisfiable. We say that $\mathcal{G}$ has *agreement soundness error* ε, if for unsatisfiable ϕ, for any assignment $\varphi_A : A \to \Sigma_A$, the A-vertices are in total disagreement on at least $1-\varepsilon$ fraction of the $b \in B$.

Feige used parallel repetition together with a coding theoretic trick to achieve agreement soundness. We show a different way to achieve agreement soundness. Our construction centers around the following combinatorial lemma.

Lemma 4.3 (Combinatorial construction)**.** *For $0 < \varepsilon < 1$, for a prime power D, and n that is a power of D, there is an explicit construction of a regular graph $H = (U,V,E)$ with $|U| = n$, V-degree D, and $|V| \le n^{O(1)}$ that satisfies the following. For every partition $U_1,\dots,U_\ell$ of U into sets such that $|U_i| \le \varepsilon |U|$ for $i = 1,\dots,\ell$, the fraction of vertices $v \in V$ with more than one neighbor in any single set U_i, is at most εD^2.*

Note that the combinatorial property could be achieved by a randomized construction, or by a construction that has a V-vertex per every possible set of D neighbors in U. However, the first construction is randomized and the second—too wasteful with a size of $\approx |U|^D$. The lemma can therefore be thought of as a *derandomization* of the randomized/full constructions.

Proof of Lemma 4.3. Our graph will be the line-point incidence graph of an affine space over a finite field. Let $U = \mathbb{F}^m$ where $\mathbb{F}$ is a finite field of order $|\mathbb{F}| = D$, and m is a natural number. Let V be the set of all affine lines in $\mathbb{F}^m$. Hence, $|V| = \binom{|U|}{2}/\binom{|\mathbb{F}|}{2}$. We connect a line $v \in V$ with a point $u \in U$ if u lies in v.

Let us show this construction satisfies the desired property. Fix a partition $U_1,\dots,U_\ell$ of U into tiny sets, $|U_i| \le \varepsilon |U|$ for $i = 1,\dots,\ell$. For every $1 \le i \le \ell$, the number of V lines that have at least two neighbors in U_i is at most $\binom{|U_i|}{2}$. Thus the total number of V-vertices with more than one neighbor in a single U_i is at most

$$
\begin{aligned}
\sum_{i=1}^{\ell} \binom{|U_i|}{2} &\le \sum_{i=1}^{\ell} \frac{|U_i|^2}{2} \\
&\le \max\{\, |U_i| \mid 1 \le i \le \ell\} \cdot \sum_{i=1}^{\ell} \frac{|U_i|}{2} \\
&\le \varepsilon |U| \cdot \frac{|U|}{2} \\
&\le \varepsilon |\mathbb{F}|^2 |V| .
\end{aligned}
$$

□

We show how to take a LABEL-COVER instance with standard soundness and convert it to a LABEL-COVER instance with total disagreement soundness, by combining it with the graph from Lemma 4.3.

Lemma 4.4. *Let $D \ge 2$ be a prime power and let n be a power of D. Let $\varepsilon > 0$. From a* LABEL-COVER *instance with soundness error $\varepsilon^2 D^2$ and B-degree n, we can construct a* LABEL-COVER *instance with agreement soundness error $2\varepsilon D^2$ and B-degree D. The transformation preserves the alphabets. The size is raised to a constant power.*

Proof. Let $\mathcal{G} = (G = (A,B,E),\Sigma_A,\Sigma_B,\Phi)$ be the original LABEL-COVER. Let $H = (U,V,E_H)$ be the graph from Lemma 4.3, where n, D and ε are as given in the current lemma. Let us use U to enumerate the neighbors of a B-vertex, i. e., there is a function $E^{\leftarrow} : B \times U \to A$ that given a vertex $b \in B$ and $u \in U$, gives us the A-vertex which is the u neighbor of b.

We create a new LABEL-COVER $(G=(A,B\times V,E'),\Sigma_A,\Sigma_B,\Phi')$. The intended assignment to every vertex $a\in A$ is the same as its assignment in the original instance. The intended assignment to a vertex $\langle b,v\rangle\in B\times V$ is the same as the assignment to b in the original game. We put an edge $e'=(a,\langle b,v\rangle)$ if $E^{\leftarrow}(b,u)=a$ and $(u,v)\in E_H$. We define $\pi_{e'}\equiv\pi_{(a,b)}$.

If there is an assignment to the original instance that satisfies c fraction of its edges, then the corresponding assignment to the new instance satisfies c fraction of its edges.

Suppose there is an assignment for the new instance $\varphi_A:A\to\Sigma_A$ in which more than $2\varepsilon D^2$ fraction of the vertices in $B\times V$ do not have total disagreement.

Let us say that $b\in B$ is "good" if for more than an εD^2 of the vertices in $\{b\}\times V$ the A-vertices do not totally disagree. Note that the fraction of good $b\in B$ is at least εD^2.

Focus on a good $b\in B$. Consider the partition of U into $|\Sigma_B|$ sets, where the set corresponding to $\sigma\in\Sigma_B$ is:

$$U_\sigma=\{\,u\in U\mid a=E^{\leftarrow}(b,u)\wedge e=(a,b)\wedge\pi_e(\varphi_A(a))=\sigma\,\}\,.$$

By the goodness of b and the property of H, there must be $\sigma\in\Sigma_B$ such that $|U_\sigma|>\varepsilon\,|U|$. We call σ the "champion" for b.

We define an assignment $\varphi_B:B\to\Sigma_B$ that assigns good vertices b their champions, and other vertices b arbitrary values. The fraction of edges that φ_A,φ_B satisfy in the original instance is at least ε^2D^2. □

Next we consider a variant of LABEL-COVER that is relevant for the reduction to SET-COVER. In this variant the prover is allowed to assign each vertex ℓ values, and an agreement is interpreted as agreement on *one* of the assignments in the list.

Definition 4.5 (List total disagreement)**.** Let $(G=(A,B,E),\Sigma_A,\Sigma_B,\Phi)$ be a LABEL-COVER. Let $\ell\ge 1$. Let $\hat{\varphi}_A:A\to\binom{\Sigma_A}{\ell}$ be an assignment that assigns each A-vertex ℓ alphabet symbols. We say that the A-vertices *totally disagree* on a vertex $b\in B$ if there are no two neighbors $a_1,a_2\in A$ of b, for which there exist $\sigma_1\in\hat{\varphi}_A(a_1)$, $\sigma_2\in\hat{\varphi}_A(a_2)$, such that

$$\pi_{e_1}(\sigma_1)=\pi_{e_2}(\sigma_2)\,,$$

where $e_1=(a_1,b),e_2=(a_2,b)\in E$.

Definition 4.6 (List agreement soundness)**.** Let $(G=(A,B,E),\Sigma_A,\Sigma_B,\Phi)$ be a LABEL-COVER for deciding membership whether a Boolean formula ϕ is satisfiable. We say that G has *list-agreement soundness error* (ℓ,ε), if for unsatisfiable ϕ, for any assignment $\hat{\varphi}_A:A\to\binom{\Sigma_A}{\ell}$, the A-vertices are in total disagreement on at least $1-\varepsilon$ fraction of the $b\in B$.

If a PCP has low error ε, then even when the prover is allowed to assign each A-vertex ℓ values, the game is still sound. This is argued in the next corollary.

Lemma 4.7 (LABEL-COVER with list-agreement soundness)**.** *Let $\ell\ge 1$, $0<\varepsilon'<1$. A LABEL-COVER with agreement soundness error ε' has list-agreement soundness error $(\ell,\varepsilon'\ell^2)$.*

Proof. Assume by way of contradiction that the LABEL-COVER instance has an assignment $\hat{\varphi}_A:A\to\binom{\Sigma_A}{\ell}$ such that on more than $\varepsilon'\ell^2$ fraction of the B-vertices, the A-vertices do not totally disagree. Define an assignment $\varphi_A:A\to\Sigma_A$ by assigning every vertex $a\in A$ a symbol picked uniformly at random from the ℓ

symbols in $\hat{\varphi}_A(a)$. If a vertex $b \in B$ has two neighbors $a_1, a_2 \in A$ that agree on b under the list assignment $\hat{\varphi}_A$, then the probability that they agree on b under the assignment φ_A is at least $1/\ell^2$. Thus, under φ_A, the expected fraction of the B-vertices that have at least two neighbors that agree on them, is more than ε'. In particular, there exists an assignment to the A-vertices, such that more than ε' fraction of the B-vertices have two neighbors that agree on them. This contradicts the agreement soundness. $\square$

The following statement summarizes the above.

Corollary 4.8. *For any $\ell = \ell(n) = \operatorname{poly}\log n$, for any constant prime power $D \geq 2$ and constant $0 < \alpha < 1$, SAT on input of size n can be reduced to a LABEL-COVER instance of size $N = poly(n)$ with alphabet size $poly(n)$, where the B-degree is D, and the list-agreement soundness error is (ℓ, α).*

Proof. Our starting point is the LABEL-COVER from Theorem 2.3 with soundness error $(2\varepsilon)^{k/2}$ so $\sqrt{(2\varepsilon)^{k/2}} \leq \alpha/2(D\ell)^2$. We apply Lemma 4.4 and Lemma 4.7. $\square$

4.2 Following Feige's reduction

In the remainder, we will show how to use Corollary 4.8 to obtain the desired hardness result for SET-COVER. The reduction is along the lines of Feige's original reduction.

For the reduction we rely on a combinatorial construction by Naor, Schulman, and Srinivasan [33]. They construct a universe together with partitions of it. Each partition covers the universe, but any cover that uses at most one set out of each partition, is necessarily large. A formal statement follows.

Lemma 4.9 (Partition system). *For natural numbers m, D and $0 < \alpha < 1$, for all $u \geq (D^{O(\log D)} \log m)^{1/\alpha}$, there is an explicit construction of a universe U of size u and partitions $\mathcal{P}_1, \ldots, \mathcal{P}_m$ of U into D sets that satisfy the following: there is no cover of U with $\ell = D \ln |U| (1-\alpha)$ sets $S_{i_1}, \ldots, S_{i_\ell}$, $1 \leq i_1 < \cdots < i_\ell \leq m$, such that set S_{i_j} belongs to partition $\mathcal{P}_{i_j}$.*

We will use the contrapositive of the lemma: if U has a cover of size at most ℓ, then this cover must contain at least two sets from the same partition. The choice of parameters of interest to us is the following: m is at most polynomial in n (m will be $|\Sigma_B|$ of the projection game), D is a sufficiently large constant, and α is a small constant.

To see why $\ell = D \ln |U| (1-\alpha)$ is to be expected (this later determines the hardness factor we get), think of the following randomized construction: each element in U corresponds to a vector in $[D]^m$, specifying for each of the m partitions, to which of its D sets it belongs. Consider a uniformly random choice of such a vector. Fix any $S_{i_1}, \ldots, S_{i_\ell}$. The probability that a random element is not covered by $S_{i_1}, \ldots, S_{i_\ell}$ is $(1-1/D)^\ell \approx \mathrm{e}^{-\ell/D}$. When $\ell = D \ln |U| (1-\alpha)$, we have $\mathrm{e}^{-\ell/D} \geq 1/|U|$, and we expect one of the $|U|$ elements in U not to be covered by $S_{i_1}, \ldots, S_{i_\ell}$. The construction of "anti-universal sets" in [33] derandomizes this randomized construction. This is the mapping from our notation to the notation in [33]: $m \to n$, $D \to b$, $\ell \to k$, U is the anti-universal set.

We now describe the reduction from a LABEL-COVER $\mathcal{G}$ as in Corollary 4.8, to a SET-COVER instance $\mathcal{SC}_\mathcal{G}$.

Apply Lemma 4.9 for $m = |\Sigma_B|$ and D which is the B-degree of the LABEL-COVER. The parameter u will be determined later. Let U be the universe, and $\mathcal{P}_{\sigma_1}, \ldots, \mathcal{P}_{\sigma_m}$ be the partitions of U. We index the partitions by Σ_B symbols $\sigma_1, \ldots, \sigma_m$. The elements of the SET-COVER instance are $B \times U$. Equivalently, each vertex $b \in B$ has a copy of the universe U. Covering this universe corresponds to satisfying the edges that touch b. There are m ways to satisfy the edges that touch b—one for every possible assignment $\sigma \in \Sigma_B$ to b. The different partitions covering U correspond to those different assignments.

For every vertex $a \in A$ and an assignment $\sigma \in \Sigma_A$ to a we have a set $S_{a,\sigma}$ in the SET-COVER instance. Taking $S_{a,\sigma}$ to the cover would correspond to assigning σ to a. Notice that a cover might consist of several sets of the form $S_{a,\cdot}$ for the same $a \in A$, which is the reason we consider list agreement. The set $S_{a,\sigma}$ is a union of subsets, one for every edge $e = (a,b)$ touching a. Suppose e is the i-th edge coming into b ($1 \le i \le D$), then the subset associated with e is $\{b\} \times S$, where S is the i-th subset of the partition $\mathcal{P}_{\varphi_e(\sigma)}$.

If we have an assignment to the A-vertices, such that all of the neighbors of b agree on one value for b, then the D subsets corresponding to those neighbors and their assignments form a partition that covers b's universe. On the other hand, if one uses only sets that correspond to totally disagreeing assignments to the neighbors, then by the definition of the partitions, covering U requires $\approx \ln|U|$ times more sets. Formally, we prove the following.

Claim 4.10. *The following hold:*

- *Completeness: If all the edges in $\mathcal{G}$ can be satisfied, then $\mathcal{SC}_\mathcal{G}$ has a set cover of size $|A|$.*
- *Soundness: Let $\ell \doteq D\ln|U|(1-\alpha)$ be as in Lemma 4.9. If $\mathcal{G}$ has agreement soundness (ℓ,α), then every set cover of $\mathcal{SC}_\mathcal{G}$ is of size more than $|A|\ln|U|(1-2\alpha)$.*

Proof. Completeness follows from taking the set cover corresponding to each of the A-vertices and its satisfying assignment.

Let us prove soundness. Assume by way of contradiction that there is a set cover C of $\mathcal{SC}_\mathcal{G}$ of size at most $|A|\ln|U|(1-2\alpha)$. For every $a \in A$ let s_a be the number of sets in C of the form $S_{a,\cdot}$. Hence, $\sum_{a\in A} s_a = |C|$. For every $b \in B$ let s_b be the number of sets in C that participate in covering $\{b\} \times U$. Then, denoting the A-degree of G by D_A,

$$\sum_{b\in B} s_b = \sum_{a\in A} s_a D_A \le D_A\,|A|\ln|U|\,(1-2\alpha) = D\,|B|\ln|U|\,(1-2\alpha)\,.$$

In other words, on average over the $b \in B$, the universe $\{b\} \times U$ is covered by at most $D\ln|U|(1-2\alpha)$ sets. Therefore, by Markov's inequality, the fraction of $b \in B$ whose universe $\{b\} \times U$ is covered by at most $D\ln|U|(1-\alpha) = \ell$ sets is at least α. By the contrapositive of Lemma 4.9 and our construction, for such $b \in B$, there are two edges $e_1 = (a_1,b), e_2 = (a_2,b) \in E$ with $S_{a_1,\sigma_1}, S_{a_2,\sigma_2} \in C$ where $\pi_{e_1}(\sigma_1) = \pi_{e_2}(\sigma_2)$.

We define an assignment $\hat{\varphi}_A : A \to \binom{\Sigma_A}{\ell}$ to the A-vertices as follows. For every $a \in A$ pick ℓ different symbols $\sigma \in \Sigma_A$ from those with $S_{a,\sigma} \in C$ (add arbitrary symbols if there are not enough). As we showed, for at least α fraction of the $b \in B$, the A-vertices will not totally disagree. □

Proof of Theorem 1.2. Fix a constant $0 < \alpha < 1$ and a prime power D. For a sufficiently large $\ell' = \Theta(\log n)$, let $\mathcal{G} = (G = (A,B,E), \Sigma_A, \Sigma_B, \Phi)$ be the LABEL-COVER with list-agreement soundness (ℓ',α) obtained from Corollary 4.8. We take $u = |U| = \Theta(|B|^{1/\alpha})$, so $u \ge (D^{O(\log D)} \log|\Sigma_B|)^{1/\alpha}$ as required for Lemma 4.9. Let $\ell = D\ln u(1-\alpha) \le \ell'$. The inapproximability ratio we get for SET-COVER from Claim 4.10 is $(1-2\alpha)\ln|U|$. Let $N = |U||B|$ be the number of elements in $\mathcal{SC}_\mathcal{G}$. We have $\ln N = (1+\alpha)\ln|U|$. The inapproximability ratio is at least $(1-3\alpha)\ln N$. Note that the reduction is polynomial in $|A|$, $|\Sigma_A|$, $|B|$, $|\Sigma_B|$ and $|U|$. Hence, the reduction is polynomial in n. This proves Theorem 1.2. □

5 Applications of the Projection Games Conjecture

In this section we describe a few applications of the PGC to hardness of approximation.

5.1 The CLOSEST-VECTOR-PROBLEM

The CLOSEST-VECTOR-PROBLEM (CVP) is to find, given a basis $b_1,\dots,b_n \in \mathbb{R}^n$ and a point $x \in \mathbb{R}^n$, the closest point to x—with respect to the ℓ_2 distance—in the lattice spanned by $b_1,\dots,b_n$, i. e., in

$$\left\{ \sum_{i=1}^{n} \alpha_i b_i \,\middle|\, \alpha_1,\dots,\alpha_n \in \mathbb{Z} \right\}.$$

Lattice problems like CVP are quite natural and have been studied a lot. One of the motivations for studying them comes from cryptography, where encryption systems believed to be secure even against quantum adversaries were built assuming the worst-case hardness of approximating lattice problems. The inapproximability factors known to be useful for cryptography are as large as $\tilde{\Omega}(n)$ for constructing collision resistant hash functions and one-way functions [29], and $\Omega(n^2)$ for public-key cryptography [38], but it is unlikely that such an approximation is NP-hard, as it (and in fact any NP-hardness of approximation to within $c\sqrt{n}$ for some constant $c > 0$) would result in a collapse of the polynomial hierarchy [1]. For more details see [28, 39].

A central question is whether one can show that lattice problems are NP-hard to approximate to within some polynomial factors $\ll \sqrt{n}$. The best existing NP-hardness result for CVP is for a factor of $\exp((\log n)^{1-\alpha})$ for any constant $\alpha > 0$ (and even for certain $\alpha = o(1)$) [12]. Assuming the PGC, we can obtain hardness of approximating CVP up to polynomial factors by a reduction of Arora, Babai, Stern and Sweedyk [3]. We state the theorem as stated by Khot [23].

Theorem 5.1 (CVP Hardness)**.** *Given a* LABEL-COVER $\mathcal{G} = (G = (A,B,E),\Sigma_A,\Sigma_B,\Phi)$ *one can construct in poly*(N) *time a lattice* $\mathcal{L}$ *in* $\mathbb{R}^N$ *and a point* $x \in \mathbb{R}^N$ *where* $N = |A|\,|\Sigma_A| + |B|\,|\Sigma_B|$*, such that the following hold.*

- *Completeness: If there is an assignment to the vertices of G that satisfies all of its edges, then the distance between x and* $\mathcal{L}$ *is at most* $\sqrt{2\,|A|\,|B|}$*.*
- *Soundness: If there is no assignment to the vertices of G that satisfies even* ε *fraction of its edges, then the distance of x and* $\mathcal{L}$ *is at least* $0.1\sqrt{|A|\,|B|\,/\varepsilon}$*.*

Hence, assuming the PGC, there exists $c > 0$*, such that approximating* CLOSEST-VECTOR-PROBLEM *to within* N^c *on an N-dimensional lattice is* NP*-hard.*

5.2 Around the approximability thresholds of CSPs

Constraint Satisfaction Problems (CSP) are defined by a set of variables $v_1,\dots,v_n$, an alphabet Σ, and constraints $\varphi_1,\dots,\varphi_m$, each depending on q variables. The number $q = O(1)$ is called the *arity* of the CSP. The task is to find an assignment to the variables that maximizes the number of satisfied constraints. One obtains specific CSPs by restricting the type of constraints. Examples include MAX-3SAT, where one is given 3CNF clauses on Boolean variables, and MAX-qLIN, where one is given linear equations over GF(2).

CSPs have been studied a lot in hardness of approximation, and for many of them we know sharp approximability thresholds. In fact, assuming the Unique Games Conjecture, we know that all CSPs over constant-sized alphabets have thresholds, where they pass from admitting polynomial time algorithms to being NP-hard [35]. For specific problems like MAX-3LIN, we know even sharper results:

Theorem 5.2 (Hardness of MAX-3LIN [16, 20])**.** *Linear* LABEL-COVER *on inputs of size n and soundness/completeness error ε can be reduced to distinguishing, given a* MAX-3LIN *instance of size $N = n\,\mathrm{poly}(1/\varepsilon)$, between the case that $(1-\varepsilon')$ fraction of the equations can be satisfied, and the case where no assignment satisfies more than $(1/2+\varepsilon')$ fraction of the equations, where $\varepsilon = \mathrm{poly}(\varepsilon')$. The reduction is linear in N.*

Hence, assuming the linear PGC, approximating MAX-3LIN *to within $1/2 + 1/N^c$ for some constant $c > 0$ is* NP*-hard.*

Note that a random assignment to the variables satisfies half of the equations in expectation, and one can always find in deterministic polynomial time an assignment that satisfies at least half of the equations. The theorem says that approximating MAX-3LIN transitions from being easy to being hard within a window of ε' at $1/2$. The width ε' determines how sharp the phase transition is. Note that at $1/2+1/N^{o(1)}$ the approximation problem is (essentially) exponentially hard assuming the exponential time hypothesis and the linear PGC. This matches an approximation algorithm by Håstad [15].

Theorem 5.2 is proved by using the Hadamard code as in [20] instead of the long code as in [16]. The advantage of the reduction in [16] is that it allows one to start with (non-linear) LABEL-COVER. Its disadvantage is that it incurs a blow-up of $N = n\exp(1/\varepsilon)$. Using [16] and Theorem 2.1, the current record, not assuming the linear PGC, is $\varepsilon' = 1/(\log\log N)^{O(1)}$.

Results analogous to Theorem 5.2 hold for other CSPs as well, e. g., for MAX-3LIN over larger finite fields, for MAX-3SAT and for other problems from Håstad's paper [16].

6 Open Problems

The main open problem is to prove (or disprove) the Projection Games Conjecture.

We believe that many more hardness of approximation results could be proved based on the PGC. Here are some concrete open problems in this direction.

1. Prove a theorem similar to Theorem 5.2 for *satisfiable* instances of MAX-3SAT.

2. Prove PGC-based hardness results for large families of CSPs similar to what is known under the Unique Games Conjecture for all CSPs [35]. A significant step in this direction was recently taken by Chan [8].

3. Prove a PGC-based hardness result for approximating SHORTEST-VECTOR-PROBLEM to within polynomial factors. Note that there is a quasi-polynomial reduction from CLOSEST-VECTOR-PROBLEM to SHORTEST-VECTOR-PROBLEM [22, 17] (see survey [23]).

4. Prove a PGC-based hardness result for approximating CLIQUE to within $N/\operatorname{poly}\log N$. Note that there is a quasi-polynomial reduction from MAX-3LIN to CLIQUE [20, 24].

Another open problem is to show equivalence between the PGC and the linear PGC.

Acknowledgments

The motivation to prove the SET-COVER result came from discussions with Ran Raz. The author would also like to thank Scott Aaronson, Zach Friggstad, Ryan O'Donnell, Muli Safra and anonymous reviewers for useful comments.

References

[1] DORIT AHARONOV AND ODED REGEV: Lattice problems in $\mathsf{NP} \cap \mathsf{coNP}$. *J. ACM*, 52(5):749–765, 2005. Preliminary version in FOCS'04. [doi:10.1145/1089023.1089025] 230

[2] NOGA ALON, DANA MOSHKOVITZ, AND SHMUEL SAFRA: Algorithmic construction of sets for k-restrictions. *ACM Trans. Algorithms*, 2(2):153–177, 2006. [doi:10.1145/1150334.1150336] 222

[3] SANJEEV ARORA, LÁSZLÓ BABAI, JACQUES STERN, AND ELIZABETH SWEEDYK: The hardness of approximate optima in lattices, codes, and systems of linear equations. *J. Comput. System Sci.*, 54(2):317–331, 1997. Preliminary version in FOCS'93. [doi:10.1006/jcss.1997.1472] 224, 230

[4] SANJEEV ARORA, CARSTEN LUND, RAJEEV MOTWANI, MADHU SUDAN, AND MARIO SZEGEDY: Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998. Preliminary versions in FOCS'92 and ECCC. [doi:10.1145/278298.278306] 223

[5] SANJEEV ARORA AND SHMUEL SAFRA: Probabilistic checking of proofs: a new characterization of NP. *J. ACM*, 45(1):70–122, 1998. Preliminary version in FOCS'92. [doi:10.1145/273865.273901] 223

[6] SANJEEV ARORA AND MADHU SUDAN: Improved low-degree testing and its applications. *Combinatorica*, 23(3):365–426, 2003. Preliminary versions in STOC'97 and ECCC. [doi:10.1007/s00493-003-0025-0] 222

[7] MIHIR BELLARE, SHAFI GOLDWASSER, CARSTEN LUND, AND ALEXANDER RUSSELL: Efficient probabilistically checkable proofs and applications to approximations. In *Proc. 25th STOC*, pp. 294–304. ACM Press, 1993. Erratum in STOC'94. [doi:10.1145/167088.167174] 224

[8] SIU ON CHAN: Approximation resistance from pairwise independent subgroups. In *Proc. 45th STOC*, pp. 447–456. ACM Press, 2013. Expanded version in ECCC. [doi:10.1145/2488608.2488665] 231

[9] MOSES CHARIKAR, MOHAMMADTAGHI HAJIAGHAYI, AND HOWARD J. KARLOFF: Improved approximation algorithms for label cover problems. *Algorithmica*, 61(1):190–206, 2011. Preliminary version in ESA'09. [doi:10.1007/s00453-010-9464-3] 224

[10] MAREK CYGAN, ŁUKASZ KOWALIK, AND MATEUSZ WYKURZ: Exponential-time approximation of weighted set cover. *Inform. Process. Lett.*, 109(16):957–961, 2009. [doi:10.1016/j.ipl.2009.05.003] 222

[11] IRIT DINUR, ELDAR FISCHER, GUY KINDLER, RAN RAZ, AND SHMUEL SAFRA: PCP characterizations of NP: Toward a polynomially-small error-probability. *Comput. Complexity*, 20(3):413–504, 2011. Preliminary versions in STOC'99 and ECCC. [doi:10.1007/s00037-011-0014-4] 224

[12] IRIT DINUR, GUY KINDLER, RAN RAZ, AND SHMUEL SAFRA: Approximating CVP to within almost-polynomial factors is NP-hard. *Combinatorica*, 23(2):205–243, 2003. Preliminary versions in ECCC and FOCS'98. [doi:10.1007/s00493-003-0019-y] 230

[13] IRIT DINUR AND DAVID STEURER: Analytical approach to parallel repetition. In *Proc. 46th STOC*, pp. 624–633. ACM Press, 2014. [ACM DL]. [doi:10.1145/2591796.2591884] 222, 223, 224

[14] URIEL FEIGE: A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998. Preliminary version in STOC'96. [doi:10.1145/285055.285059] 222

[15] JOHAN HÅSTAD: On bounded occurrence constraint satisfaction. *Inform. Process. Lett.*, 74(1-2):1–6, 2000. [doi:10.1016/S0020-0190(00)00032-6] 231

[16] JOHAN HÅSTAD: Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001. Preliminary versions in STOC'97 and ECCC. [doi:10.1145/502090.502098] 223, 231

[17] ISHAY HAVIV AND ODED REGEV: Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *Theory of Computing*, 8(23):513–531, 2012. Preliminary version in STOC'07. [doi:10.4086/toc.2012.v008a023] 231

[18] RUSSELL IMPAGLIAZZO AND RAMAMOHAN PATURI: On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. Preliminary version in CCC'99. [doi:10.1006/jcss.2000.1727] 222

[19] DAVID S. JOHNSON: Approximation algorithms for combinatorial problems. *J. Comput. System Sci.*, 9(3):256–278, 1974. [doi:10.1016/S0022-0000(74)80044-9] 222

[20] SUBHASH KHOT: Improved inapproximability results for MaxClique, chromatic number and approximate graph coloring. In *Proc. 42nd FOCS*, pp. 600–609. IEEE Comp. Soc. Press, 2001. [doi:10.1109/SFCS.2001.959936] 231, 232

[21] 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] 224

[22] SUBHASH KHOT: Hardness of approximating the shortest vector problem in lattices. *J. ACM*, 52(5):789–808, 2005. Preliminary version in FOCS'04. [doi:10.1145/1089023.1089027] 231

[23] SUBHASH KHOT: Inapproximability results for computational problems on lattices. In *The LLL Algorithm*, pp. 453–473. Springer, 2010. [doi:10.1007/978-3-642-02295-1_14] 230, 231

[24] SUBHASH KHOT AND ASHOK KUMAR PONNUSWAMI: Better inapproximability results for MaxClique, chromatic number and Min-3Lin-Deletion. In *Proc. 33rd Internat. Colloq. on Automata, Languages and Programming (ICALP'06)*, pp. 226–237, 2006. [doi:10.1007/11786986_21] 232

[25] LÁSZLÓ LOVÁSZ: On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4):383–390, 1975. [doi:10.1016/0012-365X(75)90058-8] 222

[26] CARSTEN LUND AND MIHALIS YANNAKAKIS: On the hardness of approximating minimization problems. *J. ACM*, 41(5):960–981, 1994. Preliminary version in STOC'93. [doi:10.1145/185675.306789] 222

[27] PASIN MANURANGSI AND DANA MOSHKOVITZ: Improved approximation algorithms for projection games (Extended abstract). In *ESA*, pp. 683–694, 2013. Update in CoRR. [doi:10.1007/978-3-642-40450-4_58] 224

[28] DANIELE MICCIANCIO AND SHAFI GOLDWASSER: *Complexity of Lattice Problems: A cryptographic perspective*. Volume 671 of *Kluwer Ser. in Engin. and Comput. Sci.* Kluwer, 2002. 230

[29] DANIELE MICCIANCIO AND ODED REGEV: Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS'04. [doi:10.1137/S0097539705447360] 230

[30] DANA MOSHKOVITZ: The projection games conjecture and the NP-hardness of $\ln n$-approximating set-cover. In *Proc. 15th Internat. Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'12)*, pp. 276–287, 2012. Preliminary version in ECCC. [doi:10.1007/978-3-642-32512-0_24] 221, 223

[31] DANA MOSHKOVITZ AND RAN RAZ: Sub-constant error low degree test of almost-linear size. *SIAM J. Comput.*, 38(1):140–180, 2008. Preliminary version in STOC'06. [doi:10.1137/060656838] 223

[32] DANA MOSHKOVITZ AND RAN RAZ: Two-query PCP with sub-constant error. *J. ACM*, 57(5/29), 2010. Preliminary version in FOCS'08. [doi:10.1145/1754399.1754402] 223, 225

[33] MONI NAOR, LEONARD J. SCHULMAN, AND ARAVIND SRINIVASAN: Splitters and near-optimal derandomization. In *Proc. 36th FOCS*, pp. 182–191. IEEE Comp. Soc. Press, 1995. [doi:10.1109/SFCS.1995.492475] 228

[34] DAVID PELEG: Approximation algorithms for the Label-Cover$_{\mathrm{MAX}}$ and Red-Blue Set Cover problems. *J. Discrete Algorithms*, 5(1):55–64, 2007. Preliminary version in SWAT'00. [doi:10.1016/j.jda.2006.03.008] 224

[35] PRASAD RAGHAVENDRA: Optimal algorithms and inapproximability results for every CSP? In *Proc. 40th STOC*, pp. 245–254. ACM Press, 2008. [doi:10.1145/1374376.1374414] 231

[36] RAN RAZ: A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. Preliminary version in STOC'95. [doi:10.1137/S0097539795280895] 222, 223

[37] RAN RAZ AND SHMUEL SAFRA: A sub-constant error-probability low-degree test and a sub-constant error-probability PCP characterization of NP. In *Proc. 29th STOC*, pp. 475–484. ACM Press, 1997. [doi:10.1145/258533.258641] 222

[38] ODED REGEV: On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6/34):1–40, 2009. Preliminary version in STOC'05. [doi:10.1145/1568318.1568324] 230

[39] ODED REGEV: On the complexity of lattice problems with polynomial approximation factors. In *The LLL Algorithm*, pp. 475–496. Springer, 2010. [doi:10.1007/978-3-642-02295-1_15] 230

[40] PETR SLAVÍK: A tight analysis of the greedy algorithm for set cover. *J. Algorithms*, 25(2):237–254, 1997. Preliminary version in STOC'96. [doi:10.1006/jagm.1997.0887] 222

[41] ARAVIND SRINIVASAN: Improved approximations guarantees for packing and covering integer programs. *SIAM J. Comput.*, 29(2):648–670, 1999. [doi:10.1137/S0097539796314240] 222

[42] SHERMAN K. STEIN: Two combinatorial covering theorems. *J. Combin. Theory Ser. A*, 16(3):391–397, 1974. [doi:10.1016/0097-3165(74)90062-4] 222

AUTHOR

Dana Moshkovitz
Assistant professor
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA
dmoshkov@csail .mit.edu
`http://people.csail.mit.edu/dmoshkov/`

ABOUT THE AUTHOR

DANA MOSHKOVITZ graduated from the Weizmann Institute of Science in 2008, where her advisor was Ran Raz. She was co-awarded The Haim Nessyahu Prize for the best Ph. D. in mathematics in 2008 for her thesis, titled "Two Query Probabilistic Checking of Proofs with Subconstant Error." Following post-doctoral fellowships at Princeton University and the Institute for Advanced Study, Dana joined MIT in 2010. Dana's research interests include Probabilistically Checkable Proofs and hardness of approximation, randomness in computation and coding theory.