Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT

Information Processing Letters 89 (2004) 247–254

Information Processing Letters

www.elsevier.com/locate/ipl

On the hardness of approximating label-cover

Irit Dinur *, Shmuel Safra

School of Mathematics and Computer Science, Tel-Aviv University, Tel-Aviv, Israel

Received 15 December 2002; received in revised form 13 November 2003

Communicated by L.A. Hemaspaandra

Abstract

The LABEL-COVER problem, defined by S. Arora, L. Babai, J. Stern, Z. Sweedyk [Proceedings of 34th IEEE Symposium on Foundations of Computer Science, 1993, pp. 724–733], serves as a starting point for numerous hardness of approximation reductions. It is one of six 'canonical' approximation problems in the survey of Arora and Lund [Hardness of Approximations, in: Approximation Algorithms for NP-Hard Problems, PWS Publishing Company, 1996, Chapter 10]. In this paper we present a direct combinatorial reduction from low error-probability PCP [Proceedings of 31st ACM Symposium on Theory of Computing, 1999, pp. 29–40] to LABEL-COVER showing it NP-hard to approximate to within $2^{(\log n)^{1-o(1)}}$. This improves upon the best previous hardness of approximation results known for this problem.

We also consider the MINIMUM-MONOTONE-SATISFYING-ASSIGNMENT (MMSA) problem of finding a satisfying assignment to a monotone formula with the least number of 1's, introduced by M. Alekhnovich, S. Buss, S. Moran, T. Pitassi [Minimum propositional proof length is NP-hard to linearly approximate, 1998]. We define a hierarchy of approximation problems obtained by restricting the number of alternations of the monotone formula. This hierarchy turns out to be equivalent to an AND/OR scheduling hierarchy suggested by M.H. Goldwasser, R. Motwani [Lecture Notes in Comput. Sci., Vol. 1272, Springer-Verlag, 1997, pp. 307–320]. We show some hardness results for certain levels in this hierarchy, and place LABEL-COVER between levels 3 and 4. This partially answers an open problem from M.H. Goldwasser, R. Motwani regarding the precise complexity of each level in the hierarchy, and the place of LABEL-COVER in it.

Keywords: Computational complexity; Hardness of approximation; PCP; Label-cover

1. Introduction

The LABEL-COVER problem is a combinatorial graph labelling problem defined as follows. The input is a bipartite graph $G = (U, V, E)$, two sets of labels, B_1 for U and B_2 for V, and for each edge $(u, v) \in E$, a relation $\Pi_{u,v} \subseteq B_1 \times B_2$ consisting of admissible pairs of labels for that edge. A *labelling* (f_1, f_2) is a pair of functions $f_1 : U \to 2^{B_1}$, $f_2 : V \to 2^{B_2} \setminus \{\phi\}$ assigning a subset of labels to each vertex. A labelling *covers* an edge (u, v) if for every label $a_2 \in f_2(v)$ there is a label $a_1 \in f_1(u)$ such that $(a_1, a_2) \in \Pi_{u,v}$. The goal is to find a labelling that covers all edges such that the l_p norm of the vector $(|f_1(u_1)|, |f_1(u_2)|, \ldots, |f_1(u_m)|) \in \mathbb{Z}^{|U|}$ is minimized.

* Corresponding author.
E-mail addresses: iritd@eecs.berkeley.edu (I. Dinur), safra@post.tau.ac.il (S. Safra).

0020-0190/$ – see front matter
doi:10.1016/j.ipl.2003.11.007

This problem was shown (implicitly in [10] and more formally in [2]) quasi-NP-hard to approximate to within a factor of $2^{(\log n)^{1-\delta}}$ for any constant $\delta > 0$ by showing a specific two-prover one-round interactive proof protocol, which reduces to LABEL-COVER.

We prove that LABEL-COVER is NP-hard to approximate to within $2^{(\log n)1^{1-\delta}}$ where $\delta = (\log\log n)^{-c}$ for any $c < 1/2$. This improves the best previously known results achieving NP-hardness rather than quasi-NP-hardness, and obtaining a larger factor for which hardness-of-approximation is proven. Our result also immediately strengthens the results of [6,1] and shows that the following problems are NP-hard to approximate to within a factor of $2^{(\log n)^{1-1/(\log\log n)^c}}$ for any $c < 1/2$: MMSA, MINIMUM-LENGTH-FREGE-PROOF, MINIMUM-LENGTH-RESOLUTION-REFUTATION, AND/OR SCHEDULING, LINEAR-REMOVE-PART, REMOVE-PART, SEPARATE-PAIR, FULL-DISASSEMBLY, REMOVE-SET, and SEPARATE-SET.

Remark. In [2], LABEL-COVER was reduced to the CLOSEST-VECTOR problem, the NEAREST CODEWORD problem, MAX-SATISFY, MIN-UNSATISFY, learning half-spaces in the presence of errors, and a number of other problems. Unfortunately, their reduction, is not from general LABEL-COVER, but rather relies on a special additional property of the LABEL-COVER instance that they construct. Namely that the relations associated with each edge are partial functions: every label for u can be covered by at most one label for v. This property is inherently missing in our reduction, and indeed hardness results for the aforementioned problems seem to require more work than is present in our direct reduction.

A formula-depth hierarchy

We also consider a related problem called MINIMUM-MONOTONE-SATISFYING-ASSIGNMENT (MMSA) that was defined in [1], and shown there to be as hard as LABEL-COVER. Given a formula φ over a monotone basis, the problem is to find a satisfying assignment for φ with a minimum number of 1's. This problem was considered in [1] since it reduces to the problem of finding the length of a propositional proof, a problem of considerable interest in proof-theory. Subsequently, Umans [13] obtained an n^{ε} hardness result for a related problem: that of finding the minimum weight assignment for a circuit whose set of accepting strings is monotone. In that problem the circuit itself is not necessarily given by a formula that is written over a monotone basis, so the hardness result does not carry over to our case.

We show that the MMSA problem can be viewed as a generalization of the LABEL-COVER problem. We examine a hierarchy of approximation problems formed by restricting the depth of the monotone formula in the MMSA problem. This hierarchy is equivalent to a hierarchy of AND/OR scheduling pointed out in [6]. A monotone formula is said to be of depth i if it has $i-1$ alternations between AND and OR. A depth-i formula is called Π_i (Σ_i) if the first level of alternation is an AND (OR). It is easy to see that the complexity of MMSA restricted to Σ_{i+1} formulae is equivalent the complexity of MMSA restricted to Π_i formulae, denoted MMSA_i.

Each MMSA_i is at least as hard to approximate as MMSA_{i-1}. MMSA_1 is trivially solvable in polynomial time. MMSA_2, is already quite harder, and actually a simple approximation-preserving reduction from SET-COVER to MMSA_2 was shown in [1], implying that MMSA_2 is NP-hard to approximate to within logarithmic factors [12]. In fact, the two problems can be easily shown to be equivalent, thus the same greedy algorithm for SET-COVER [7,9] approximates MMSA_2 to within a factor of $\ln n$. We know of no previous hardness result for MMSA_3. A reduction from LABEL-COVER to MMSA_4 was shown independently in [1] and [6].

We show how to translate MMSA_3 to LABEL-COVER, altogether placing LABEL-COVER somewhere between levels 3 and 4 in this hierarchy. This partially answers an open question from [6] of whether or not LABEL-COVER is equivalent to level 4 in the hierarchy. Furthermore, we examine the (previously unknown) hardness of MMSA_3 and via a reduction from PCP to MMSA_3 show that it is NP-hard to approximate to within the above large factors. This immediately carries over for MMSA_i for every $i \geqslant 3$ and for LABEL-COVER. Our reductions all involve a polynomial sized blow-up, thus the hardness-of-approximation ratios are polynomially related. For the asymptotic approximation ratios discussed here, this polynomial blow-up is irrelevant.

Table 1

Formula depth	Approximation algorithm	NP-hardness factor
MMSA_1	1	–
MMSA_2	$\ln n$	$\Omega(\log n)$
$\mathrm{MMSA}_{\geqslant 3}$	n	$2^{\log^{1-\mathrm{o}(1)} n}$

If we denote the relation *reducible with a polynomially related approximation-ratio* by $\ll$ we can write:

$$\mathrm{PCP} \ll \mathrm{MMSA}_3 \ll \text{LABEL-COVER} \ll \mathrm{MMSA}_4 \ll \cdots \ll \mathrm{MMSA}_i.$$

We summarize the above in Table 1.

Technique

We show a direct reduction to LABEL-COVER from low error-probability PCP with parameters D and ε. Namely, we begin with a gap-SAT instance consisting of Boolean constraints. These constraints each depend on D variables, and the variables take values in $\{1, \ldots, \lceil 1/\varepsilon \rceil\}$. The PCP theorem states that it is NP-hard to distinguish between the 'yes' case where all of the constraints are satisfiable, and the 'no' case where every assignment satisfies no more than an ε fraction of the constraints. The focus of [4] was on $D = \mathrm{O}(1)$, and thus only an error-probability of $\varepsilon = 2^{-(\log n)^{1-\delta}}$ for any constant $\delta > 0$ was claimed. This alone strengthens the hardness of LABEL-COVER from quasi-NP-hardness to NP-hardness, but with the same hardness-factor as before. For our purposes however, the best result is obtained by choosing $D = \log\log^c n$ for any $c < 1/2$ and $\varepsilon = 2^{-\log^{1-1/\mathrm{O}(D)} n}$. These parameters give the result claimed above. Notice that our direct reduction immediately implies that a stronger PCP characterization of NP—e.g., one with a polynomially-small error-probability and constant depend as conjectured in [3]—would immediately give NP-hardness for approximating LABEL-COVER to within n^c for some constant $c > 0$.

Structure of the paper

Our main result for LABEL-COVER is proven in Section 2. The hardness result for MMSA_3 is proven in Section 3, via a reduction from PCP. We then show, in Section 4 a reduction from MMSA_3 to LABEL-COVER thus placing LABEL-COVER between levels 3 and 4 in the 'MMSA' hierarchy. This also re-establishes the hardness result for LABEL-COVER already shown in Section 2.

2. Label-cover

The LABEL-COVER problem is defined as follows.

Definition 1 (LABEL-COVER (LC_p)). The input to the label-cover problem is a bipartite graph $G = (U, V, E)$, two sets of labels, B_1 for U and B_2 for V, and for each edge $(u, v) \in E$, a relation $\Pi_{u,v} \subseteq B_1 \times B_2$ consisting of admissible pairs of labels for that edge.

A *labelling* (f_1, f_2) is a pair of functions $f_1 : U \to 2^{B_1}$, $f_2 : V \to 2^{B_2} \setminus \{\phi\}$ assigning a subset of labels to each vertex. The l_p-cost of the labelling is the l_p norm of the vector $(|f_1(u_1)|, |f_1(u_2)|, \ldots, |f_1(u_m)|) \in \mathbb{Z}^{|U|}$. A labelling *covers* an edge (u, v) if for every label $a_2 \in f_2(v)$ there is a label $a_1 \in f_1(u)$ such that $(a_1, a_2) \in \Pi_{u,v}$. A *total-cover* of G is a labelling that covers every edge. The problem LC_p is to find a total-cover with minimal l_p-cost ($1 \leqslant p \leqslant \infty$).

In this section we show a direct reduction from PCP to LABEL-COVER with l_p norm, $1 \leqslant p \leqslant \infty$, such that the approximation factor is preserved.

Let us denote $g_c(n) \stackrel{def}{=} 2^{\log^{1-1/\log\log^c n} n}$. Our reduction will imply that LABEL-COVER is NP-hard to approximate to within factor $g_c(n)$ for any $c < 1/2$. Our starting point is the PCP theorem from [4].

Theorem 1 (PCP Theorem [4]). *Let $c < 1/2$ be arbitrary and let $D \leqslant \log\log^c n$. Let $\Psi = \{\psi_1, \ldots, \psi_n\}$ be a system of boolean constraints over variables $X = \{x_1, \ldots, x_{n'}\}$ such that each boolean constraint depends on D variables, and each variable takes values in $\mathcal{F}$ where $|\mathcal{F}| = \mathrm{O}(2^{(\log n)^{1-1/\mathrm{O}(D)}})$. It is* NP*-hard to distinguish between the following two cases*:

Yes: There is an assignment to the variables such that all $\psi_1, \ldots, \psi_n$ are satisfied.

No: No assignment can satisfy more than $\mathrm{O}(1)/|\mathcal{F}|$ *fraction of the ψ_i's.*

The following is the main theorem in this section.

Theorem 2. *For any $c < 1/2$, and any $1 \leqslant p \leqslant \infty$, LABEL-COVER_p is NP-hard to approximate to within a factor of $g = g_c(n)$.*

Proof. The proof follows by reduction from PCP. Choose some $c < c' < 1/2$, let $\mathcal{F}$ be such that $|\mathcal{F}| = \Theta(g_{c'}(n))$, and let $\Psi = \{\psi_1, \dots, \psi_n\}$ be a PCP instance as in the above Theorem 1. For a constraint $\psi \in \Psi$ and a variable $x \in X$, we write $x \in \psi$ when ψ depends on x, and denote $\Psi_x \stackrel{def}{=} \{\psi \in \Psi \mid x \in \psi\}$.

We construct from Ψ a bipartite graph $G = (U, V, E)$ as follows. Let $U \stackrel{def}{=} \{u_1, \dots, u_{nD}\}$ have a vertex for every appearance of a variable in Ψ, and let $V \stackrel{def}{=} [n]$ have a vertex for every constraint $\psi \in \Psi$. We denote $U(x) \subset U$ the set of vertices corresponding to the variable x. A vertex $j \in V$ is connected to *all* appearances of all of the variables in ψ_j. Formally,

$$E \stackrel{def}{=} \{(u, j) \mid u \in U(x) \text{ and } x \in \psi_j\}.$$

We set $B_1 \stackrel{def}{=} \mathcal{F}$ and $B_2 \stackrel{def}{=} \mathcal{F}^D$. For an edge $(u, j) \in E$, assume $u \in U(x)$ and x is the ith variable in ψ_j, and define

$$\Pi_{u,j} = \{(a_i, (a_1, \dots, a_D)) \mid \psi_j(a_1, \dots, a_D) = \mathtt{True}\}.$$

Proposition 1 (Completeness). *If there is a satisfying assignment for Ψ, then there is a total-cover for G with l_∞ cost 1, and l_1 cost $n \cdot D$.*

Proof. Let $\mathcal{A}: X \to \mathcal{F}$ be an assignment satisfying all of Ψ. Define for each $u \in U(x)$, $f_1(u) \stackrel{def}{=} \{\mathcal{A}(x)\}$ and for each $j \in V$, $f_2(j) \stackrel{def}{=} \{(\mathcal{A}(x_{i_1}), \dots, \mathcal{A}(x_{i_D})) \mid \psi_j$ depends on $x_{i_1}, \dots, x_{i_D}\}$ (these are both singleton sets). This is a total-cover of l_∞ cost 1 and l_1-cost $n \cdot D$. $\square$

We next show that if Ψ is a 'no' instance, then any label-cover has l_∞ cost more than g. This is formulated in a contrapositive manner as follows.

Proposition 2 (Soundness$_\infty$). *If there is a total-cover for G with l_∞ cost g, then there is an assignment $\mathcal{A}$ satisfying a $g^{-D} > \mathrm{O}(1)/|\mathcal{F}|$ fraction of Ψ (so Ψ is not a 'no' instance).*

Proof. Let (f_1, f_2) be a labelling for G that is a total-cover with l_∞-cost g, i.e.,

$$\max_i(|f_1(v_i)|) = g.$$

We define a random assignment $\mathcal{A}$ for the variables X by choosing for every variable x_i a value uniformly at random from $f_1(u)$ where $u \in U(x_i)$ is an arbitrary vertex in $U(x_i)$. Since the labelling is a total-cover, each label $r \in f_2(v_j)$ corresponds to an assignment that satisfies ψ_j and such that $r|_{x_i} \in f_1(u)$ for every vertex $u \in U(x_i)$ and variable x_i appearing in ψ_j. Thus, a constraint ψ_j is satisfied with probability $|f_2(v_j)|/g^D \geqslant g^{-D}$, so the expected fraction of constraints satisfied by $\mathcal{A}$ is also $\geqslant g^{-D}$. There must be an assignment that attains the expectation, and satisfies at least a g^{-D} fraction of the constraints in Ψ.

Note that for the $g = g_c(n)$ chosen above $g^{-D} > \mathrm{O}(1)/|\mathcal{F}|$ because $|\mathcal{F}| = \mathrm{O}(g_{c'}(n))$ for $c' > c$, thus Ψ is not a 'no' instance. $\square$

We next show that if Ψ is a 'no' instance, then any label-cover has l_1 cost more than g. This again, is formulated in a contrapositive manner as follows.

Proposition 3 (Soundness$_1$). *If there is a total-cover for G with l_1-cost $g \cdot nD$, then there is an assignment $\mathcal{A}$ satisfying $\geqslant (1/2) \cdot (1/(2D \cdot g)^D) > \mathrm{O}(1)/|\mathcal{F}|$ fraction of Ψ.*

Proof. Let (f_1, f_2) be a total-cover with l_1 cost $g \cdot nD$. For every variable x, define $A(x) \stackrel{def}{=} \bigcap_{u \in U(x)} f_1(u) \subseteq \mathcal{F}$ (this set is nonempty since (f_1, f_2) is a total cover). Recall that $\Psi_x \subseteq \Psi$ denotes the set of constraints that depend on x. If $u \in U(x)$ then $|A(x)| \leqslant |f_1(u)|$, hence

$$\sum_{i=1}^{n'} |\Psi_{x_i}| \cdot |A(x_i)| \leqslant \sum_{u \in U} |f_1(u)| = g \cdot nD. \qquad (*)$$

Consider the random procedure of choosing a constraint $\psi \in_R \Psi$ uniformly at random and then choosing a variable $x \in_R \psi$ uniformly at random. The probability of choosing x is $|\Psi_x|/(nD)$. Eq. $(*)$ is equivalent to $E(|A(x)|) \leqslant g$ where $E(|A(x)|)$ denotes the expectation of $|A(x)|$ for x chosen by the above random procedure.

We call a variable x for which $|A(x)| > 2D \cdot g$, a *bad* variable. By the Markov inequality

$$\Pr_x[|f_1(x)| > 2D \cdot E(|A(x)|)] < \frac{1}{2D}$$

which means that the probability of hitting a bad variable is less than $1/(2D)$.

$$\begin{aligned}\frac{1}{2D} &\geqslant \Pr_{\psi\in\Psi_x,\, x\in\psi}[x \text{ is bad}]\\ &= \Pr_{\psi\in_R\Psi}[\psi \text{ contains a bad variable}]\\ &\quad\times \Pr_{x\in\psi}[x \text{ is bad } \mid \psi \text{ contains a bad variable}]\\ &\geqslant \Pr_{\psi\in_R\Psi}[\psi \text{ contains a bad variable}]\cdot\frac{1}{D}.\end{aligned}$$

Multiplying by D, we deduce that at least half of the constraints $\psi \in_R \Psi$ contain no bad variable. Next, define a random assignment $\mathcal{A}$ for Ψ by choosing, for every variable x, a random value $a \in A(x)$, $\mathcal{A}(x) \stackrel{def}{=} a$. For a constraint ψ_i and a value $r \in f_2(v_i)$, the probability that each variable $x \in \psi_i$ was assigned $a = r|_x$ is $\prod_{x\in\psi_i}(1/|A(x)|)$ (recall that r satisfies ψ_i so this is a lower bound on the probability that ψ_i is satisfied by $\mathcal{A}$). For constraints that contain no bad variable, this probability is $\geqslant 1/(2D\cdot g)^D$. Hence the expected fraction of constraints (of those containing no bad variable) that are satisfied by $\mathcal{A}$ is $\geqslant 1/(2D\cdot g)^D$. Thus, there exists an assignment $\mathcal{A}$ that attains this expectation, i.e., that satisfies a $\geqslant 1/(2D\cdot g)^D$ fraction of the constraints that contain no bad variables. Thus $\mathcal{A}$ satisfies a $\geqslant 1/2(2D\cdot g)^D$ fraction of all of the constraints.

Let us see that for the above chosen $g = g_c(n)$, $1/(2\cdot(2D\cdot g)^D) > \mathrm{O}(1)/|\mathcal{F}|$. Indeed taking inverses and considering the $\log\log$ of both sides,

$$\begin{aligned}&\log\log 2\cdot(2D\cdot g)^D\\ &\quad\leqslant \mathrm{O}(\log D) + \log\log g\\ &\quad= \mathrm{O}(\log\log\log n) + \left(1 - \frac{1}{(\log\log n)^c}\right)\log\log n\\ &\quad< \log\log\bigl(g_{c'}(n)\bigr)\\ &\quad= \log\log|F|.\end{aligned}$$

Thus Ψ is not a 'no' instance. $\square$

Propositions 1 and 3 imply that distinguishing between the case where there is a total-cover for G whose l_1 cost is nD or $g\cdot nD$ would enable distinguishing between 'yes' and 'no' PCP instances, hence it is NP-hard. Similarly, Propositions 1 and 2 imply the same about distinguishing between the case where there is a total-cover for G whose l_∞ cost is 1 or g. As for other l_p norms $1 < p < \infty$, the same reduction gives an inapproximability gap of $\sqrt[p]{g}$ for *any* l_p norm with $1 < p < \infty$. This follows since in the completeness case a 'yes' PCP instance translates into a zero-one vector, whose l_p norm is $\sqrt[p]{nD}$. In the soundness case, since for $x \in \mathbb{Z}^n$, $\|x\|_1 = \sum|x_i| \leqslant \sum|x_i|^p = \|x\|_p^p$, any total-cover whose l_p norm is at most $\sqrt[p]{gnD}$ has l_1 norm at most gnD and the soundness argument for the l_1 case carries through. The inapproximability gap thus obtained is $\sqrt[p]{gnD}/\sqrt[p]{nD} = \sqrt[p]{g}$, which, for constant p, is of the same order of magnitude as before. $\square$

3. Reducing PCP to MMSA$_3$

The MINIMUM-MONOTONE-SATISFYING-ASSIGNMENT (MMSA) problem is defined as follows,

Definition 2 (*MMSA*). Given a monotone formula $\varphi(x_1,\ldots,x_k)$ over the basis $\{\wedge,\vee\}$, find a satisfying assignment $A:\{x_1,\ldots,x_k\}\to\{0,1\}$ (i.e., such that $\varphi(A(x_1),\ldots,A(x_k)) = \mathtt{True}$), minimizing the *weight* $\sum_{i=1}^k A(x_i)$.

MMSA$_i$ is the restriction of MMSA to formulae of depth i. For example, MMSA$_3$ is the problem of finding a minimal-weight assignment for a formula written as an AND of ORs of ANDs.

In this section we show a direct reduction from PCP to MMSA$_3$, that preserves the approximation factor.

Theorem 3. *For any $c < 1/2$, it is* NP-*hard to approximate* MMSA$_3$ *to within $g_c(n) \stackrel{def}{=} 2^{\log^{1-1/\log\log^c n} n}$.*

Proof. Again, our starting point is the low error-probability PCP theorem, Theorem 1. Fix $g = g_c(n)$, and fix $c < c' < 1/2$ arbitrarily. Take $\mathcal{F}$ to be such that $|\mathcal{F}| = \Theta(g_{c'}(n))$, and $D = \mathrm{O}(\log\log^{c'} n)$. Let Ψ be a PCP instance as in Theorem 1. For a fixed $\psi \in \Psi$, we denote the set of satisfying assignments for it $R_\psi \subseteq \mathcal{F}^D$. For an assignment $r \in R_\psi$ and a variable $x \in \psi$ we write $r|_x \in \mathcal{F}$ to denote the restriction of r to x.

We construct the monotone formula Φ over the following set of literals

$$T \stackrel{def}{=} \bigcup_{x\in X}\bigl\{T[x,\psi,a] \mid \psi\in\Psi_x,\ a\in\mathcal{F}\bigr\}.$$

This set has cardinality $nD \cdot |\mathcal{F}|$. The pair of variable x and assignment a for it will be represented by the conjunction $L[x,a] \stackrel{def}{=} \bigwedge_{\psi \in \Psi_x} T[x,\psi,a]$ that can be read as "a is assigned to x". We define the formula $\Phi(T)$ by

$$\Phi(T) \stackrel{def}{=} \bigwedge_{\psi\in\Psi} \bigvee_{r\in R_\psi} \bigwedge_{x\in\psi} L[x, r|_x].$$

This is a depth-3 formula, since the conjunction of conjunctions is still a conjunction.

Proposition 4 (Completeness). *If Ψ is satisfiable, then there is a satisfying assignment for Φ, whose weight is $n \cdot D$.*

Proof. Let $\mathcal{A}: X \to \mathcal{F}$ be a satisfying assignment for Ψ. Define an assignment $\mathcal{A}': T \to \{\mathtt{True}, \mathtt{False}\}$ for the literals of Φ by setting $\mathcal{A}'(T[x,\psi,a]) = \mathtt{True}$ iff $\mathcal{A}(x) = a$. This assignment clearly satisfies Φ, and has weight exactly nD. □

Proposition 5 (Soundness). *If there is a satisfying assignment for Φ whose weight is gnD, then there is an assignment satisfying a $1/(2(2Dg)^D)$ fraction of Ψ.*

The proof of this proposition is very similar to the proof of Proposition 3.

Proof. Let $\mathcal{A}_\Phi: T \to \{\mathtt{True}, \mathtt{False}\}$ be a weight-gnD satisfying assignment for Φ. For each variable $x \in X$, let $A(x) \stackrel{def}{=} \{a \in \mathcal{F} \mid \mathcal{A}_\Phi(L[x,a]) = \mathtt{True}\}$. $A(x)$ is nonempty since x appears in some constraint ψ, and for each $\psi \in \Psi$ there must be some r for which $\bigwedge_{x\in\psi} L[x, r|_x] = \mathtt{True}$ because $\mathcal{A}_\Phi$ satisfies Φ.

$L[x,a]$ contains $|\Psi_x|$ literals that, if $a \in A(x)$, are by definition set to $\mathtt{True}$. These are distinct for distinct x's, thus

$$\sum_{x\in X} |\Psi_x| \cdot \bigl|A(x)\bigr| \leqslant g \cdot nD.$$

Consider the procedure of choosing a constraint $\psi \in_R \Psi$ uniformly at random and then choosing a variable $x \in_R \psi$ uniformly at random. The probability of choosing x is $|\Psi_x|/(nD)$. The above equation is thus equivalent to $E(|A(x)|) \leqslant g$ where $E(|A(x)|)$ denotes the expectation of $|A(x)|$ where x is chosen by the above procedure.

We call a variable x for which $|A(x)| > 2D \cdot g$, a *bad* variable. The Markov inequality yields

$$\Pr_x\bigl[\bigl|A(x)\bigr| > 2D \cdot E\bigl(\bigl|A(x)\bigr|\bigr)\bigr] < \frac{1}{2D}$$

which means that the probability of hitting a bad variable is less than $1/(2D)$.

$$\begin{aligned}\frac{1}{2D} &\geqslant \Pr_{\psi\in\Psi, x\in\psi}[x \text{ is bad}] \\ &= \Pr_{\psi\in_R\Psi}[\psi \text{ contains a bad variable}] \\ &\quad\times \Pr_{x\in\psi}[x \text{ is bad} \mid \psi \text{ contains a bad variable}] \\ &\geqslant \Pr_{\psi\in_R\Psi}[\psi \text{ contains a bad variable}] \cdot \frac{1}{D}.\end{aligned}$$

Multiplying by D, we deduce that at least half of the constraints $\psi \in_R \Psi$ contain no bad variable.

Next, we define a random assignment $\mathcal{A}$ for Ψ by choosing, for every variable x, a random value $a \in A(x)$, $\mathcal{A}(x) \stackrel{def}{=} a$. For each constraint $\psi \in \Psi$ there is at least one value $r \in R_\Psi$ with $\bigwedge_{x\in\psi} \mathcal{A}_\Phi(L[x, r|_x]) = \mathtt{True}$ since $\mathcal{A}_\Phi$ satisfies Φ. The probability that each variable $x \in \psi$ was assigned $a = r|_x \in A(x)$ is $\prod_{x\in\psi}(1/|A(x)|)$. For constraints that contain no bad variable, this probability is $\geqslant 1/(2D \cdot g)^D$. Hence there is an assignment that satisfies at least

$$\frac{1}{2} \cdot \frac{1}{(2D \cdot g)^D}$$

fraction of the constraints.

Since $1/(2(2Dg)^D) > \mathrm{O}(1)/|\mathcal{F}|$, we deduce that Ψ is not a 'no' PCP instance. □

We saw in Proposition 4 that if Ψ is a PCP 'yes' instance then there is a weight-nD satisfying assignment for Φ. On the other hand, if Ψ is a PCP 'no' instance (i.e., any assignment satisfies no more than a $\mathrm{O}(1)/|\mathcal{F}|$ fraction of the constraints), then there cannot be even a weight-gnD satisfying assignment for Φ. Otherwise Proposition 5 would imply that there is an assignment satisfying a $1/2 \cdot (2Dg)^D > \mathrm{O}(1)/|\mathcal{F}|$ fraction of the constraints (the last inequality holds because $c' > c$). Thus, distinguishing between the case where the monotone formula has a satisfying assignment of weight nD or gnD is NP-hard because it enables distinguishing between 'yes' and 'no' PCP instances. This completes the proof of the theorem. □

4. Reducing MMSA_3 to LABEL-COVER

In this section we show a reduction from MMSA_3 to LABEL-COVER. This shows that MMSA_3 is no harder than LABEL-COVER, and (together with the reduction from [1]) places LABEL-COVER between level 3 and 4 in the 'MMSA-hierarchy'. It also re-establishes the result in Section 2 showing NP-hardness for approximating LABEL-COVER.

An instance of MMSA_3 is a formula

$$\Phi \stackrel{def}{=} \bigwedge_{i=1}^{I} \bigvee_{j=1}^{J} \bigwedge_{k=1}^{K} T_{i,j,k}$$

where the $T_{i,j,k}$ are literals from the set $\{\mathtt{x}_1, \ldots, \mathtt{x}_L\}$ for some $L \leqslant I \cdot J \cdot K$ (by repeating literals we may assume wlog that all conjunctions are of the same size, and similarly all disjunctions). We construct a bipartite graph $G = (U, V, E)$ with vertices $U \stackrel{def}{=} \{u_1, \ldots, u_L\}$ for the literals, and

$$V \stackrel{def}{=} \{v_{i,w} \mid 1 \leqslant w \leqslant W, 1 \leqslant i \leqslant I\}$$

for W copies of the I disjunctions (where W is chosen large enough, say $W = L$). The edges in E connect every literal to the disjunctions in which it appears,

$$E \stackrel{def}{=} \big\{(u_l, v_{i,w}) \mid \exists j, k, T_{i,j,k} = \mathtt{x}_l\big\}.$$

The sets of possible labels are $B_1 \stackrel{def}{=} \{0, 1, \ldots, W\}$ for U and $B_2 \stackrel{def}{=} \{1, \ldots, J\}$ for V.

For $j = 1, \ldots, J$, denote

$$T_{i,j} = \{T_{i,j,k} \mid 1 \leqslant k \leqslant K\}.$$

If a vertex $v = v_{i,w}$ is labelled by j, we differentiate between two kinds of neighbors u_l of v: those with $\mathtt{x}_l \in T_{i,j}$ and those with $x_l \notin T_{i,j}$. For an edge $e = (u_l, v_{i,w})$, we construct the relation Π_e so that the two kinds of neighbors are 'covered' differently,

$$\Pi_e \stackrel{def}{=} \big\{(w, j) \mid \mathtt{x}_l \in T_{i,j}\big\} \cup \big\{(0, j) \mid \mathtt{x}_l \notin T_{i,j}\big\}.$$

Note that for every label j for v there is at least one u_l for which $\mathtt{x}_l \in T_{i,j}$, thus labelling $u_1, \ldots, u_L$ with 0 cannot be a total-cover.

Proposition 6 (Completeness). *If there is a satisfying assignment for Φ whose weight is t, then there is a total-cover for G with l_1 cost $L + t \cdot W = (t + 1) \cdot W$.*

Proof. Let $\mathcal{A}$ be a weight-t satisfying assignment for Φ. Define a cover as follows. For every $u_l \in U$ set

$$f_1(u_l) \stackrel{def}{=} \begin{cases} \{0, 1, \ldots, W\} & \mathcal{A}(\mathtt{x}_l) = \mathtt{True}, \\ \{0\} & \text{otherwise}. \end{cases}$$

For every $v_{i,w} \in V$ let $f_2(v_{i,w}) \stackrel{def}{=} \{j^*\}$ where j^* is the smallest index for which

$$\bigwedge_{k=1}^{K} \mathcal{A}(T_{i,j^*,k}) = \mathtt{True}$$

(such an index j^* exists because $\mathcal{A}$ satisfies Φ). Obviously f_1, f_2 are nonempty, and the l_1 cost of the labelling is exactly $L + t \cdot W$.

Let us show that the labelling (f_1, f_2) is a total cover. Let $e = (u_l, v_{i,w})$ be an arbitrary edge, and let $j \in f_2(v_{i,w})$. By definition of f_2, j is such that $\mathcal{A}(\mathtt{x}_l) = \mathtt{True}$ for all $\mathtt{x}_l \in T_{i,j}$. Thus, for an index l with $\mathtt{x}_l \in T_{i,j}$, by definition $f_1(u_l) = \{0, 1, \ldots, W\}$ and e is covered by (w, j). If $\mathtt{x}_l \notin T_{i,j}$ then $(0, j) \in \Pi_e$ so e is covered because $0 \in f_1(u_l)$. □

Proposition 7 (Soundness). *If there is a total-cover for G with l_1 cost tW, then there is a satisfying assignment for Φ whose weight is t.*

Proof. Let (f_1, f_2) be a total cover with l_1 cost tW. Since $\forall u \in U$ $f_1(u) \subseteq \{0, 1, \ldots, W\}$, and

$$\sum_{u \in U} |f_1(u)| = tW,$$

there must be at least one $w^* > 0$ for which $|\{u \mid w^* \in f_1(u)\}| \leqslant t$. We claim that the assignment $\mathcal{A}$ defined by assigning $\mathtt{x}_l$ the value $\mathtt{True}$ if and only if $w^* \in f_1(u_l)$, satisfies Φ. Note that $\mathcal{A}$'s weight cannot exceed t.

Fix i. We will show that the ith disjunction is satisfied. Consider the vertex v_{i,w^*} and a label $j \in f_2(v_{i,w^*}) \neq \emptyset$. As before, define $T_{i,j} = \{T_{i,j,k} \mid k = 1, \ldots, K\}$. We will show that the jth conjunction of the ith disjunction is satisfied (thus satisfying the whole disjunction). For this purpose we need to show that every literal $\mathtt{x}_l \in T_{i,j}$ is assigned $\mathtt{True}$, or in other words $w^* \in f_1(u_l)$. But this is immediate since there is no other way of covering the edges $e \stackrel{def}{=} (u_l, v_{i,w^*})$, and (f_1, f_2) is a total-cover. □

Summarizing Propositions 6 and 7, we see that if the original formula Φ had a satisfying assignment of weight t_0, then the LABEL-COVER instance has a total-cover whose l_1-cost is $W(t_0+1)$. If, on the other hand, every satisfying assignment for Φ has weight $> t = g_c(n) \cdot t_0$, then every total-cover has l_1-cost $> g_c(n) \cdot Wt_0$.

In the previous section we saw that it is NP-hard to distinguish between the case where the minimum weight assignment is at most t_0 or at least $g_c(n) \cdot t_0$, thus it is NP-hard to approximate LABEL-COVER to within a factor of $g_c(n)Wt_0/W(t_0+1) \geqslant g_c(n)/2 = \Omega(2^{(\log n)^{1-1/D}})$ where $D = (\log\log n)^c$ for any $c < 1/2$. The proof for other l_p norms follows, as before, since $1 \leqslant p < \infty$ norms approximate each other for integer-valued vectors.

5. Discussion and open questions

*A depend-*2 *PCP characterization of* NP

In [2] LABEL-COVER was used to prove the hardness of the CLOSEST-VECTOR problem along with several other problems. However, they used a slightly modified version of LABEL-COVER, in which the relation Π_e for each edge is actually a *function* from B_1 to B_2. In our result, Π_e is a function from B_2 to B_1 and inherently cannot be extended to this version. This obstacle could be overcome had we known a low error-probability PCP characterization of NP with *exactly two* provers (i.e., a PCP constraint-system where each constraint accesses exactly two variables, called depend-2-PCP). Compare this to the known low error-probability PCP characterization of NP [12,4] where each constraint depends on a constant (> 2) number of variables. Whether or not such a characterization exists remains an open question. Note that it is highly unlikely that this problem is in P since such an interactive proof protocol for NP exists [8,5, 11], with a quasi-polynomial blow-up.

The MMSA hierarchy

We considered a hierarchy of approximation problems, equivalent to that in [6]. We showed a new hardness-of-approximation result for it (starting from the third level). Are higher levels in this hierarchy even harder to approximate, perhaps to within some polynomial n^{ε} factor? Such a result would immediately strengthen the known hardness results for the aforementioned problems in [6,1].

We know that LABEL-COVER resides between levels 3 and 4 in this hierarchy. However, the inapproximability factor of LABEL-COVER is the same as that of MMSA_i for $i \geqslant 3$. Is this an indication that the hierarchy collapses, or is there really a difference in the hardness of hierarchy levels for $i \geqslant 3$?

References

[1] M. Alekhnovich, S. Buss, S. Moran, T. Pitassi, Minimum propositional proof length is NP-hard to linearly approximate, manuscript, 1998.

[2] S. Arora, L. Babai, J. Stern, Z. Sweedyk, The hardness of approximate optima in lattices, codes and linear equations, in: Proceedings of 34th IEEE Symposium on Foundations of Computer Science, 1993, pp. 724–733.

[3] M. Bellare, S. Goldwasser, C. Lund, A. Russell, Efficient multi-prover interactive proofs with applications to approximation problems, in: Proceedings of 25th ACM Symposium on Theory of Computing, 1993, pp. 113–131.

[4] I. Dinur, E. Fischer, G. Kindler, R. Raz, S. Safra, PCP characterizations of NP: Towards a polynomially-small error-probability, in: Proceedings of 31st ACM Symposium on Theory of Computing, 1999, pp. 29–40.

[5] U. Feige, L. Lovász, Two-prover one-round proof systems: Their power and their problems, in: Proceedings of 24th ACM Symposium on Theory of Computing, 1992, pp. 733–741.

[6] M.H. Goldwasser, R. Motwani, Intractability of assembly sequencing: Unit disks in the plane, in: Proceedings of the Workshop on Algorithms and Data Structures, in: Lecture Notes in Comput. Sci., vol. 1272, Springer-Verlag, Berlin, 1997, pp. 307–320.

[7] D.S. Johnson, Approximation algorithms for combinatorial problems, J. Comput. System Sci. 9 (1974) 256–278.

[8] D. Lapidot, A. Shamir, Fully parallelized multi prover protocols for NEXPTIME, in: Proceedings of 32nd IEEE Symposium on Foundations of Computer Science, 1991, pp. 13–18.

[9] L. Lovász, On the ratio of optimal integral and fractional covers, Discrete Math. 13 (1975) 383–390.

[10] C. Lund, M. Yannakakis, On the hardness of approximating minimization problems, J. ACM 41 (5) (1994) 960–981.

[11] R. Raz, A parallel repetition theorem, SIAM J. Comput. 27 (3) (1998) 763–803.

[12] R. Raz, S. Safra, A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP, in: Proceedings of 29th ACM Symposium on Theory of Computing, 1997, pp. 475–484.

[13] C. Umans, Hardness of approximating Σ_2^p minimization problems, in: 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS '99), 1999, pp. 465–474.