On the Hardness of Approximating Label Cover

Irit Dinur* Shmuel Safra*

November 1, 2002

Abstract

The Label-Cover problem, defined in [ABSS93], 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 [AL97]. In this paper we present a direct combinatorial reduction from low error-probability PCP [DFK+99] 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 least number of 1s, [ABMP98]. We define a hierarchy of approximation problems obtained by restricting the number of alternations the monotone formula. This hierarchy turns out to be equivalent to an AND/OR scheduling hierarchy suggested in [GM97]. 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 [GM97] regarding the precise complexity of each level in the hierarchy, and the place of Label-Cover in it.

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 $f_1 : U \to 2^{B_1}, f_2 : V \to 2^{B_2}$ 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.

This problem was shown (implicitly in [LY94] and more formally in [ABSS93]) quasi-NP-hard to approximate to within a factor of $2^{\log^{1-\delta} n}$ 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 $2^{\log^{1-\delta} n}$ where $\delta = \log\log^{-c} n$ 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 [GM97, ABMP98] and shows that the following problems are NP-hard to approximate to within a factor of $2^{\log^{1-1/\log\log^c n} n}$ 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.

* School of Mathematical Sciences, Tel Aviv University, ISRAEL

Remark. In [ABSS93], 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 in our direct reduction.

A Formula-Depth Hierarchy

We also consider a related problem called MINIMUM-MONOTONE-SATISFYING-ASSIGNMENT (MMSA) that was defined in [ABMP98], and shown there to be as hard as LABEL-COVER. Given a monotone formula φ the problem is to find a satisfying assignment for φ with a minimum number of 1's. This problem was considered in [ABMP98] since it reduces to the problem of finding the length of a propositional proof, a problem of considerable interest in proof-theory. Although our LABEL-COVER result strengthens the hardness for MMSA, we note that [Uma99] subsequently obtained an even better $n^{1-\varepsilon}$ hardness result for this problem without going through a reduction from LABEL-COVER.

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 [GM97]. 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} formulas is equivalent the complexity of MMSA restricted to Π_i formulas, 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 [ABMP98], implying that MMSA_2 is NP-hard to approximate to within logarithmic factors [RS97]. In fact, the two problems can be easily shown to be equivalent, thus the same greedy algorithm for SET-COVER [Joh74, Lov75] 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 [ABMP98] and [GM97].

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 [GM97] 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 follows through for MMSA_i (for every $i \geq 3$) 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.

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

$$\text{PCP} \ll \text{MMSA}_3 \ll \text{LABEL-COVER} \ll \text{MMSA}_4 \ll \ldots \ll \text{MMSA}_i$$

We summarize the above in the following table:

Formula Depth	Approximation Algorithm	NP-Hardness Factor
MMSA_1	1	–
MMSA_2	$\ln n$	$\Omega(\log n)$
$\text{MMSA}_{\geq 3}$	n	$2^{\log^{1-o(1)} n}$

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 functions. These Boolean functions each depend on D variables, and the variables range over $\{1 \ldots 1/\varepsilon\}$. The PCP theorem states that it is NP-hard to distinguish between the 'yes' case where the whole system is satisfiable, and the 'no' case where every assignment satisfies no more than an ε fraction of the local-tests. The focus of [DFK+99] was on $D = O(1)$, and thus only an error-probability of $\varepsilon = 2^{-\log^{1-\delta} n}$ 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 $\varepsilon = 2^{-\log^{1-1/O(D)} n}$. 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 [BGLR93] – 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 **($\mathbf{LC}_p$)**) *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 $f_1 : U \to 2^{B_1}, f_2 : V \to 2^{B_2}$ 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 ($1 \leq p \leq \infty$).*

In this section we show a direct reduction from PCP to Label-Cover with l_p norm, $1 \leq p \leq \infty$, 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 [DFK+99],

Theorem 1 (PCP Theorem [DFK+99]) *Let $c < 1/2$ be arbitrary and let $D \leq \log\log^c n$. let $\Psi = \{\psi_1, ..., \psi_n\}$ be a system of Boolean functions over variables $X = \{x_1, .., x_{n'}\}$ such that each Boolean function depends on D variables, and each variable ranges over $\mathcal{F}$ where $|\mathcal{F}| = O(2^{(\log n)^{1-1/O(D)}})$. It is NP-hard to distinguish between the following two cases:*

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

No: No assignment can satisfy more than $\frac{1}{|\mathcal{F}|}$ fraction of the ψ_i's.

In this section we prove LABEL-COVER to be NP-hard to approximate to within a factor of g, where $g = g_c(n)$ is fixed for some arbitrary $c < 1/2$.

Theorem 2 *For any $c < \frac{1}{2}$, and any $1 \leq p \leq \infty$, 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 $|\mathcal{F}| = O(g_{c'}(n))$, and let $\Psi = \{\psi_1, \ldots, \psi_n\}$ be a PCP instance as in the above theorem. For a test $\psi \in \Psi$ 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 $G = (U, V, E)$ with $U \stackrel{def}{=} \{u_1, ..., u_{nD}\}$ consisting of a vertex for every appearance of a variable in Ψ and $V \stackrel{def}{=} [n]$ consisting of a vertex for every test $\psi \in \Psi$. 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 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

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

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 $f_2(v_j) \stackrel{def}{=} \{(\mathcal{A}(x_{i_1}), .., \mathcal{A}(x_{i_D})) \mid \psi_j \text{ depends on } x_{i_1}, .., x_{i_D}\}$ (these are both singleton sets). This is a total-cover of l_∞ cost 1 and l_1-cost $n \cdot D$. ■

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 $g^{-D} > \frac{1}{|\mathcal{F}|}$ fraction of Ψ (and 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 $U(x_i)$. 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 test ψ_j is satisfied with probability $|f_2(v_j)|/g^D \geq g^{-D}$, so the expected number of tests satisfied by $\mathcal{A}$ is also $\geq g^{-D}$. There must be an assignment that attains the expectation, and satisfies at least g^{-D} fraction of the tests in Ψ.

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

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 $\geq \frac{1}{2} \cdot \frac{1}{(2D \cdot g)^D} > \frac{1}{|\mathcal{F}|}$ 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 non-empty since (f_1, f_2) is a total cover). Recall $\Psi_x \subseteq \Psi$ denoted the set of tests that depend on x. If $u \in U(x)$ then $|A(x)| \leq |f_1(u)|$, hence

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

Consider the random procedure of choosing a test $\psi \in_R \Psi$ uniformly at random and then choosing a variable $x \in_R \psi$ uniformly at random. The probability of choosing x is $\frac{|\Psi_x|}{nD}$. Equation $(*)$ is equivalent to $E(|A(x)|) \leq g$ where $E(|A(x)|)$ denotes the expectation of $|A(x)|$ for x is 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

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

which means that the probability of hitting a bad variable is less than $\frac{1}{2D}$.

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

Multiplying by D, we deduce that at least half of the tests $\psi \in_R \Psi$ 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 test 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 $\Pi_{x \in \psi_i} \frac{1}{|A(x)|}$ (recall that r satisfies ψ_i so this is a lower bound on the probability that ψ_i is satisfied by $\mathcal{A}$). For tests that contain no bad variable, this probability is $\geq \frac{1}{(2D\cdot g)^D}$. Hence the expected fraction of tests (of those containing no bad variable) that are satisfied by $\mathcal{A}$ is $\geq \frac{1}{(2D\cdot g)^D}$. Thus, there exists an assignment $\mathcal{A}$ that attains this expectation, i.e. that satisfies $\geq \frac{1}{(2D\cdot g)^D}$ fraction of the tests that contain no bad variables. Thus $\mathcal{A}$ satisfies a $\geq \frac{1}{2(2D\cdot g)^D}$ fraction of all of the tests.

Note that for the above chosen $g = g_c(n)$, $\frac{1}{2\cdot(2D\cdot g)^D} > 1/|\mathcal{F}|$, thus Ψ is not a 'no' instance. ■

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. The above can easily be generalized for *any* l_p norm, $1 \leq p \leq \infty$. ■

3 Reducing PCP to MMSA$_3$

The Minimum-Monotone-Satisfying-Assignment (MMSA) problem is defined as follows,

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

MMSA$_i$ is the restriction of MMSA to formulas 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 $c < \frac{1}{2}$, 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 $|\mathcal{F}| = O(g_{c'}(n))$, and $D = O(\log\log^{c'} n)$. 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

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

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 $\mathcal{A}' : T \to \{\texttt{True}, \texttt{False}\}$ for the literals of Φ by setting $\mathcal{A}'(T[x,\psi,a]) = \texttt{True}$ iff $\mathcal{A}(x) = a$. This assignment clearly satisfies Φ, and has weight exactly nD. ■

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

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

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

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

$$\sum_{x\in X} |\Psi_x| \cdot |A(x)| \le g \cdot nD\,.$$

Consider the procedure of choosing a test $\psi \in_R \Psi$ uniformly at random and then choosing a variable $x \in_R \psi$ uniformly at random. The probability of choosing x is $\frac{|\Psi_x|}{nD}$. The above equation is thus equivalent to $E(|A(x)|) \le g$ 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

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

which means that the probability of hitting a bad variable is less than $\frac{1}{2D}$.

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

Multiplying by D, we deduce that at least half of the tests $\psi \in_R \Psi$ 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 test $\psi \in \Psi$ there is at least one value $r \in R_\Psi$ with $\bigwedge_{x\in\psi} \mathcal{A}_\Phi(L[x,r|_x]) = \texttt{True}$ since $\mathcal{A}_\Phi$ satisfies Φ. The probability that each variable $x \in \psi$ was assigned $a = r|_x \in A(x)$ is $\Pi_{x\in\psi}\frac{1}{|A(x)|}$. For tests that contain no bad variable, this probability is $\ge \frac{1}{(2D\cdot g)^D}$. Hence there is an assignment that satisfies at least

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

fraction of the tests.

Since $\frac{1}{2(2Dg)^D} > \frac{1}{|\mathcal{F}|}$, 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 was a PCP 'no' instance (i.e. any assignment satisfies no more than $1/|\mathcal{F}|$ fraction of the tests), then there cannot be even a weight-gnD satisfying assignment for Φ. Otherwise Proposition 5 would imply that there is an assignment satisfying $1/2 \cdot (2Dg)^D > 1/|\mathcal{F}|$ fraction of the tests (the last inequality follows mainly 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 [ABMP98]) 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 to within the same factor.

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 $\{\mathtt{x}_1, .., \mathtt{x}_L\}$ for some $L \le I \cdot J \cdot K$ (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 $U \stackrel{def}{=} \{u_1, .., u_L\}$ for the literals, and $V \stackrel{def}{=} \bigcup_{w=1}^{W} \{v_{1,w}, .., v_{I,1}\}$ 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,

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

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

For $j = 1, ..., J$, denote $T_{i,j} = \{T_{i,j,k} \mid 1 \le k \le K\}$. 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,

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

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 $u_1, .., u_L$ with 0 cannot be a total-cover.

Proposition 6 (Completeness) *If there is a satisfying assignment for Φ with weight 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

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

For every $v_{i,w} \in V$ let $f_2(v_{i,w}) \stackrel{def}{=} \{j^*\}$ where j^* is the smallest index for which $\bigwedge_{k=1}^{K} \mathcal{A}(T_{i,j^*,k}) =$ `True` (such an index j^* exists because $\mathcal{A}$ satisfies Φ). Obviously f_1, f_2 are non-empty, 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 $\mathcal{A}(\mathtt{x}_l) =$ `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 $f_1(u_l) = \{0, 1, ..., W\}$ 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 $g \cdot tW$, then there is a satisfying assignment for Φ with weight gt.*

Proof: Let (f_1, f_2) be a total cover with l_1 cost $gt \cdot W$. Since $\forall u \in U \quad f_1(u) \subseteq \{0, 1, .., W\}$, and $\sum_{u \in U} |f_1(u)| = gt \cdot W$, there must be at least one $w^* > 0$ for which $|\{u \mid w^* \in f_1(u)\}| \leq gt$. We claim that the assignment $\mathcal{A}$ (whose weight cannot exceed gt) defined by assigning $\mathtt{x}_l$ the value `True` if and only if $w^* \in f_1(u_l)$, satisfies Φ.

Fix i. We will show that the ith disjunction is satisfied. Consider the vertex v_{i,w^*} and a label $j \in f_2(v_{i,w^*}) \neq \phi$. As before, define $T_{i,j} = \{T_{i,j,k} \mid k = 1, .., K\}$. 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 `True`, or in other words $w^* \in f_1(u_l)$. 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, then the Label-Cover instance has a total-cover whose l_1-cost is $W(t+1)$. If, on the other hand, every satisfying assignment for Φ has weight $> gt$, then every total-cover has l_1-cost $> g \cdot tW$. This completes the reduction.

Choosing $g = g_c(n)$ and by the result in the previous section we deduce that it is NP-hard to approximate Label-Cover to within a factor of $\frac{gtW}{W(t+1)} \geq g/2 = \Omega(2^{\log^{1-1/D} n})$ where $D = \log\log^c n$ for any $c < 1/2$. The proof for other l_p norms follows similarly.

5 Discussion and Open Questions

A Depend-2 PCP Characterization of NP

In [ABSS93] 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 test-system where each tests accesses exactly two variables, called depend-2-PCP). Compare this to the known low error-probability PCP characterization of NP [RS97, DFK+99] where each test 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 [LS91, FL92, Raz98], with a quasi-polynomial blow-up.

The MMSA Hierarchy

We considered a hierarchy of approximation problems, equivalent to that in [GM97]. 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 [GM97, ABMP98].

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

References

[ABMP98] M. Alekhnovich, S. Buss, S. Moran, and T. Pitassi. Minimum propositional proof length is NP-hard to linearly approximate. Manuscript, 1998.

[ABSS93] S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes and linear equations. In *Proc. 34th IEEE Symp. on Foundations of Computer Science*, pages 724–733, 1993.

[AL97] S. Arora and C. Lund. Hardness of approximations. In D. Hochbaum, editor, *Approximation Algorithms for NP-Hard Problems.* PWS Publishing, 1997.

[BGLR93] M. Bellare, S. Goldwasser, C. Lund, and A. Russell. Efficient multi-prover interactive proofs with applications to approximation problems. In *Proc. 25th ACM Symp. on Theory of Computing*, pages 113–131, 1993.

[DFK+99] I. Dinur, E. Fischer, G. Kindler, R. Raz, and S. Safra. PCP characterizations of NP: Towards a polynomially-small error-probability. In *Proc. 31st ACM Symp. on Theory of Computing*, 1999.

[FL92] U. Feige and L. Lovász. Two-prover one-round proof systems: Their power and their problems. In *Proc. 24th ACM Symp. on Theory of Computing*, pages 733–741, 1992.

[GM97] Michael H. Goldwasser and Rajeev Motwani. Intractability of assembly sequencing: Unit disks in the plane. In *Proceedings of the Workshop on Algorithms and Data Structures*, volume 1272 of *Lecture Notes in Computer Science*, pages 307–320. Springer-Verlag, 1997.

[Joh74] D. S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.

[Lov75] L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.

[LS91] D. Lapidot and A. Shamir. Fully parallelized multi prover protocols for NEXPTIME. In *Proc. 32nd IEEE Symp. on Foundations of Computer Science*, pages 13–18, 1991.

[LY94] Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998.

[RS97] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proc. 29th ACM Symp. on Theory of Computing*, pages 475–484, 1997.

[Uma99] C. Umans. Hardness of approximating $\sum_2^p$ minimization problems. In *40th Annual Symposium on Foundations of Computer Science (FOCS '99)*, pages 465–474, Washington - Brussels - Tokyo, October 1999. IEEE.