Electronic Colloquium on Computational Complexity, Report No. 15 (1999)

On the Hardness of Approximating Label Cover

Irit Dinur* Shmuel Safra*

April 23, 1999

Abstract

The LABEL-COVER problem was introduced in [ABSS93] and shown there to be quasi-NP-hard to approximate to within a factor of $2^{\log^{1-\delta} n}$ for any *constant* $\delta > 0$. This combinatorial graph problem has been utilized [ABSS93, GM97, ABMP98] for showing hardness-of-approximation of numerous problems. We present a direct combinatorial reduction from low error-probability $\mathcal{PCP}$ [DFK+] to LABEL-COVER. This improves on [ABSS93] in two ways. First, it improves the previous hardness-of-approximation factor known for LABEL-COVER, achieving a factor of $2^{\log^{1-\delta} n}$ where $\delta = 1/\log\log^c n$ for any constant $c < 1/2$. Furthermore, we show approximating LABEL-COVER is NP-hard for these large factors, compared to the *quasi* NP-hardness, obtained previously.

Our result for LABEL-COVER immediately strengthens the known hardness result for several approximation problems as mentioned above, for example the MINIMUM-MONOTONE-SATISFYING-ASSIGNMENT (MMSA) problem [ABMP98]. Furthermore, we examine a hierarchy of approximation problems obtained by restricting the depth of the monotone formula in MMSA. 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.

Introduction

The LABEL-COVER problem, implicit in [LY94], was first formally defined in [ABSS93]. The input to the LABEL-COVER problem (see definition 1 for a full description) is a bipartite graph $G = (U, V, E)$, a set of possible labels for each vertex, and a relation for each edge consisting of admissible pairs of labels for that edge. A *labeling* is an assignment of a subset of labels to each vertex. A labeling *covers* an edge (u, v) if for every label assigned to v there is a label assigned to u such that together they make an admissible pair according to the above relation. The goal is to find a minimal labeling that covers all of the edges.

This problem was shown [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. In [Hoc97] the LABEL-COVER problem is presented as one of six 'canonical' problems for proving hardness-of-approximation. Indeed, [ABSS93] reduced LABEL-COVER 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. Their reduction, however, is not from general LABEL-COVER, rather

* School of Mathematical Sciences, Tel Aviv University, ISRAEL

ISSN 1433-8092

relies on a special additional property[1] of the LABEL-COVER instance that they construct. While this property is inherently missing in our reduction, we note that the hardness of the CLOSEST-VECTOR problem has independently been strengthened [DKS98, DKRS99], and from this follows the hardness for the additional problems mentioned above.

Goldwasser and Motwani [GM97] showed an approximation preserving reduction from LABEL-COVER to an AND/OR SCHEDULING problem, thus implying the hardness of approximating to within the above factors a number of assembly sequencing problems: LINEAR-REMOVE-PART, REMOVE-PART, SEPARATE-PAIR, FULL-DISASSEMBLY, REMOVE-SET, AND SEPARATE-SET.

Another related problem called MINIMUM-MONOTONE-SATISFYING-ASSIGNMENT (MMSA) was defined in [ABMP98]. 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 shown [ABMP98] to reduce to the problem of finding the length of a propositional proof, a problem of considerable interest in proof-theory. Given some natural proof-system (a specific set of axioms and some inference rules), and given a tautology φ, the problem is to find the length of its shortest proof (measured in steps or symbols). A proof is defined to be a series of tautologies each being either an instanciation of an axiom or a result of applying an inference rule on previously obtained tautologies. Approximating this problem was shown [ABMP98] to be as hard as approximating MMSA (for almost all natural proof-systems). They showed an A-reduction (i.e. a reduction that preserves hardness-of-approximation ratio to within a constant factor, see definition in [KST97]) from MMSA to MINIMUM-LENGTH-FREGE-PROOF and to MINIMUM-LENGTH-RESOLUTION-REFUTATION, where the length is measured either by symbols or by steps.

A direct reduction from LABEL-COVER to MMSA has been shown in [ABMP98]. Combining our result with this reduction implies that a stronger $\mathcal{PCP}$ characterization of NP – e.g. one with a polynomially-small error-probability – would immediately give NP-hardness for approximating MMSA to within n^c for some constant $c > 0$. The question of whether there is a polynomial approximation (within *any* polynomial factor) for propositional proof lengths is of great interest in proof-theory, and a negative reply would imply that a polynomial-time automatic prover does not exist.

Our Results

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.

A Formula-Depth Hierarchy

In addition, 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

[1] namely that after fixing an edge, the relation is a partial function

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 the 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 $\mathcal{PCP}$ 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:

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

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 $\mathcal{PCP}$ with parameters D and ε. Namely, we begin with a system of local-tests (Boolean functions) each depending on D variables ranging over $\{1 \ldots 1/\varepsilon\}$. The $\mathcal{PCP}$ 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+] 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 $\mathcal{PCP}$ 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 1. The hardness result for MMSA_3 is proven in section 2, via a reduction from $\mathcal{PCP}$. We then show, in section 3 a reduction from MMSA_3 to LABEL-COVER thus placing LABEL-COVER between levels 3 and 4 in the hierarchy. This re-establishes the hardness result for LABEL-COVER already shown in section 1. We choose not to omit section 1 since it shows a direct simple proof for the hardness of LABEL-COVER.

1 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 $G = (V_1, V_2, E)$ (where $E \subseteq V_1 \times V_2$), and a set of possible labels B_1 and B_2 for vertices V_1 and V_2 respectively. Also included in the input is a relation $\Pi \subseteq E \times B_1 \times B_2$ that consists of admissible pairs of labels for each edge. A* labeling *of the graph is a pair of functions $(\mathcal{P}_1, \mathcal{P}_2)$ where $\mathcal{P}_i : V_i \to 2^{B_i}$ for $i = 1, 2$; in other words, $(\mathcal{P}_1, \mathcal{P}_2)$ assigns a set of labels – to each vertex of the graph. The l_p-*cost *of the labeling is the l_p norm of the vector $(|\mathcal{P}_1(v_1^1)|, \ldots, |\mathcal{P}_1(v_1^n)|) \in \mathbf{Z}^{|V_1|}$ (where $V_1 = \{v_1^1, \ldots, v_1^n\}$). A labeling is said to* cover *an edge $e = (v_1, v_2)$ if both $\mathcal{P}_1(v_1)$ and $\mathcal{P}_2(v_2)$ are non-empty, and for every label $b_2 \in \mathcal{P}_2(v_2)$ there exists a $b_1 \in \mathcal{P}_1(v_1)$ such that $(e, b_1, b_2) \in \Pi$. A* total-cover *of G is a labeling 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 $\mathcal{PCP}$ 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 $\mathcal{PCP}$ theorem from [DFK+],

Theorem 1 (PCP Theorem [DFK+]) *Let $\Psi = \{\psi_1, ..., \psi_n\}$ be a system of local-tests over variables $V = \{x_1, .., x_{n'}\}$ such that each local-test depends on $D = \log\log^c n$ variables (for any $c < 1/2$), and each variable ranges over a field $\mathcal{F}$ 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 $\frac{2}{|\mathcal{F}|}$ fraction of the ψ_i's.

We will show LABEL-COVER to be NP-hard to approximate to within a factor of g, where $g = g_c(n)$ is fixed for some arbitrary $c < 1/2$. Choose some $c < c' < 1/2$, let $\mathcal{F}$ be a field with $|\mathcal{F}| = O(g_{c'}(n))$, and let $\Psi = \{\psi_1, \ldots, \psi_n\}$ be a $\mathcal{PCP}$ instance as in the above theorem.

We construct from Ψ a graph $G = (U, V, E)$ with $U \stackrel{def}{=} \{u_1, \ldots, u_{n \cdot D}\}$ having a vertex for every appearance of a variable in Ψ and $V \stackrel{def}{=} \{v_1, .., v_n\}$ having a vertex for every test $\psi \in \Psi$.

We denote $U(x) \subset U$ the set of vertices corresponding to a variable x. Every test is connected by an edge to *every* appearance of each of its variables. The edges will be

$$E \stackrel{def}{=} \{(u, v_j) \mid \varphi_j \text{ depends on a variable } x \text{ and } u \in U(x)\}$$

The possible set of labels for U is $B_1 \stackrel{def}{=} \mathcal{F}$ the range of the variables, and for V is $B_2 \stackrel{def}{=} \mathcal{F}^D$ the set of possible assignments to the tests. The relation Π will consist of all $((u, v_j), r, r')$ where the restriction of $r' \in \mathcal{F}^D$ to the variable x for which $u \in U(x)$ gives r:

$$\Pi \stackrel{def}{=} \{((u, v_j), r|_x, r) \mid 1 \le j \le n,\ u \in U(x),\ \psi_j \text{ depends on } x,\ r \text{ satisfies } \psi_j\}$$

Proposition 1 (Soundness) *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 $\mathcal{A} : V \to \mathcal{F}$ be a satisfying assignment for Ψ. Define $\mathcal{P}_1(v_i) \stackrel{def}{=} \{\mathcal{A}(x_i)\}$ and $\mathcal{P}_2(d_j) \stackrel{def}{=} \{(\mathcal{A}(x_{i_1}), .., \mathcal{A}(x_{i_D})) \mid \psi_j \text{ depends on } x_{i_1}, .., x_{i_D}\}$. Obviously, this is a total-cover of l_∞ cost 1 and l_1-cost $n \cdot D$. ■

Proposition 2 (Completeness$_\infty$) *If there is a total-cover for G with l_∞-cost g, then there is an assignment $\mathcal{A}$ satisfying $1/g^D$ fraction of Ψ. Hence if $g \le |\mathcal{F}|^{\frac{1}{2D}}$, then Ψ is satisfiable.*

(note that this indeed holds for the $g = g_c(n)$ chosen above).

Proof: Let $(\mathcal{P}_1, \mathcal{P}_2)$ be a labeling for G that is a total-cover with l_∞-cost g, i.e.

$$\max(|\mathcal{P}_1(v_i)|) = g$$

We define a random assignment for the variables V by choosing for every variable x_i a value uniformly at random from $\mathcal{P}_1(u)$ where $u \in U(x_i)$ is arbitrary, say the vertex in $U(x_i)$ with minimal index. Every label in $\mathcal{P}_2(v_j)$ must satisfy ψ_j, otherwise it cannot be covered by the definition of Π. A test ψ_j is satisfied with probability $|\mathcal{P}_2(v_j)|/g^D \ge 1/g^D$ since each value $r \in \mathcal{P}_2(v_j)$ corresponds to an assignment that satisfies ψ_j and such that $r|_{x_i} \in \mathcal{P}_1(u)$ for every vertex $u \in U(x_i)$ and variable x_i appearing in ψ_j. We deduce the existence of an assignment that satisfies at least $1/g^D = \frac{1}{\sqrt{|\mathcal{F}|}}$ fraction of the tests in Ψ. ■

Proposition 3 (Completeness$_1$) *If there is a total-cover for G with l_1-cost $g \cdot nD$, then there is an assignment $\mathcal{A}$ satisfying $> \frac{1}{2} \cdot \frac{1}{(2D \cdot g)^D}$ fraction of Ψ. Hence if $g \le \frac{|\mathcal{F}|^{\frac{1}{2D}}}{2D}$, then Ψ is satisfiable.*

(note that this indeed holds for the $g = g_c(n)$ chosen above).

Proof: Let $(\mathcal{P}_1, \mathcal{P}_2)$ be a total-cover with l_1 cost $g \cdot nD$. For every variable x, define $A(x) \stackrel{def}{=} \bigcap_{u \in U(x)} \mathcal{P}_1(u) \subseteq \mathcal{F}$. Denoting by $\Psi_{x_i} \subseteq \Psi$ the set of tests that depend on the variable x_i, it follows that

$$\sum_{i=1}^{n'} |\Psi_{x_i}| \cdot |A(x_i)| \le \sum_{u \in U} |\mathcal{P}_1(u)| = g \cdot nD$$

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 equivalent to $E(|A(x)|) \leq g$ where $E(|A(x)|)$ denotes the expectation of $|A(x)|$ for 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

$$\Pr_x [|\mathcal{P}_1(x)| > a \cdot E(|A(x)|)] < \frac{1}{a}$$

which means that (substituting $a \stackrel{def}{=} 2D$) 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.

Finally, we apply the 'weak notion of satisfiability' paradigm as follows. Define a random assignment $\mathcal{A}$ for Ψ by choosing, for every variable x, a random value $a \in A(x)$, $A_R(x) \stackrel{def}{=} a$. For a test ψ_i and a value $r \in \mathcal{P}_2(d_i)$, 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 A_R). For tests that contain no bad variable, this probability is $\geq \frac{1}{(2D \cdot g)^D}$. Hence if $g \leq \frac{|\mathcal{F}|^{\frac{1}{2D}}}{2D}$ there is an assignment that satisfies at least

$$\frac{1}{2} \cdot \frac{1}{(2D \cdot g)^D} = \frac{1}{2 \cdot \sqrt{|\mathcal{F}|}}$$

fraction of the tests, which implies that Ψ is satisfiable. ■

Remark. It is easy to see that the above carries over for *any* l_p norm, $1 \leq p \leq \infty$.

2 Reducing $\mathcal{PCP}$ 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 $A : \{x_1, .., x_k\} \to \{0, 1\}$, $\varphi(A(x_1), .., A(x_k)) = \texttt{TRUE}$ minimizing $\sum_{i=1}^k A(x_i)$.*

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 $\mathcal{PCP}$ to MMSA$_3$, that preserves the approximation factor. Let us denote $g_c(n) \stackrel{def}{=} 2^{\log^{1-1/\log\log^c n} n}$. Our reduction will imply that MMSA$_3$ is NP-hard to approximate to within factor $g_c(n)$ for any $c < 1/2$. As before, our starting point is theorem 1 ($\mathcal{PCP}$).

Fix $g = g_c(n)$ for some arbitrary $c < 1/2$, and fix $c < c' < 1/2$ arbitrarily. Take $\mathcal{F}$ to be a field with $|\mathcal{F}| = O(g_{c'}(n))$, and $D = O(\log\log^{c'} n)$. Let Ψ be a $\mathcal{PCP}$ instance as in theorem 1. For a test $\psi \in \Psi$ and a variable $x \in V$, we write $x \in \psi$ when ψ depends on x, and denote $\Psi_x \stackrel{def}{=} \{\psi \in \Psi \mid x \in \psi\}$. We also denote the set of satisfying assignments for $\psi \in \Psi$ by $R_\psi \subseteq \mathcal{F}^D$.

We construct the monotone formula Φ over the literals B as follows.

$$B \stackrel{def}{=} \bigcup_{x \in V} \left\{ B[x^\psi, a] \mid \psi \in \Psi_x,\ a \in \mathcal{F} \right\}$$

This set has one literal $B[x^\psi, a]$ for every variable $x \in V$, and test $\psi \in \Psi_x$ in which it appears, and value for it $a \in \mathcal{F}$ (altogether $|B| = nD \cdot |\mathcal{F}|$). The pair of variable x and assignment a for it will be represented by the conjunction $L[x, a] \stackrel{def}{=} \bigwedge_{\psi \in \Psi_x} B[x^\psi, a]$. We set

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

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 $\mathcal{A} : V \to \mathcal{F}$ be a satisfying assignment for Ψ. Define an assignment $\mathcal{A}' : B \to$ $\{\mathtt{true}, \mathtt{false}\}$ for the literals of Φ by setting $\mathcal{A}'(B[x^\psi, a]) = \mathtt{true}$ iff $\mathcal{A}(x) = a$. This is a weight-nD assignment, and obviously satisfies Φ. ■

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

Proof: Let $\mathcal{A}_\Phi : B \to \{\mathtt{true}, \mathtt{false}\}$ be a weight-gnD satisfying assignment for Φ. For each variable $x \in V$, let $A(x) \stackrel{def}{=} \{a \in \mathcal{F} \mid \mathcal{A}_\Phi(L[x, a]) = \mathtt{true}\}$. Since $\mathcal{A}_\Phi$ satisfies Φ and since every variable x appears in some test $x \in \psi$, $A(x) \neq \phi$. It follows that

$$\sum_{x \in V} |\Psi_x| \cdot |A(x)| \leq g \cdot nD$$

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 $E(|A(x)|) \leq g$ where $E(|A(x)|)$ denotes the expectation of $|A(x)|$ with x chosen by the above procedure.

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

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

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

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

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

Finally, we apply the 'weak notion of satisfiability' paradigm as follows. 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 $\bigwedge_{x\in\psi} \mathcal{A}_\Phi(L[x, r|_x]) = \mathtt{true}$ since $\mathcal{A}_\Phi$ satisfies Φ. The probability that each variable $x \in \psi$ was assigned $a = r|_x \in A(x)$ is $\Pi_{x\in\psi}\frac{1}{|A(x)|}$. For tests that contain no bad variable, this probability is $\geq \frac{1}{(2D\cdot g)^D}$. Hence there is an assignment that satisfies at least

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

fraction of the tests. ■

We saw in proposition 4 that if Ψ were a $\mathcal{PCP}$ 'yes' instance then there is a weight-nD satisfying assignment for Φ. On the other hand, if Ψ was a $\mathcal{PCP}$ 'no' instance (i.e. any assignment satisfies no more than $2/|\mathcal{F}|$ 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 $1/2\cdot(2Dg)^D > 1/g^{2D} > 2/|\mathcal{F}|$ fraction of the tests (the last inequality follows mainly because $c' > c$).

3 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 1 showing NP-hardness for approximating Label-Cover to within the above factor.

An instance of MMSA$_3$ is a formula

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

where the $B_{i,j,k}$ are literals from the set $\{\mathrm{x}_1,..,\mathrm{x}_L\}$ for some $L \leq I\cdot J\cdot K$. We construct a graph $G = (V_1, V_2, E)$ with vertices $V_1 \stackrel{def}{=} \{v_1,..,v_L\}$ for the literals, and $V_2 \stackrel{def}{=} \bigcup_{w=1}^{W}\{d_1^w,..,d_I^w\}$ for W copies of the I disjunctions (where W is chosen large enough, e.g. $W = L$). The edges in E connect every literal to the disjunctions in which it appears,

$$E \stackrel{def}{=} \{(v_l, d_i^w)\mid \exists j,k,\ B_{i,j,k} = \mathrm{x}_l\}$$

The sets of possible labels are $B_1 \stackrel{def}{=} \{0,1,\ldots,W\}$ and $B_2 \stackrel{def}{=} \{1,\ldots,J\cdot W\}$. The relation Π is defined by setting for every edge $e = (v_l, d_i^w)$ and label $jw \in B_2$ an element $(e, w, jw) \in \Pi$ if the literal x_l appears in the j^{th} conjunction of the i^{th} disjunction (i.e. $\exists k,\ B_{i,j,k} = \mathrm{x}_l$). If not, we set $(e, 0, jw) \in \Pi$.

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 $v_l \in V_1$ set

$$\mathcal{P}_1(v_l) \stackrel{def}{=} \begin{cases} \{0,1,..,W\} & \mathcal{A}(\mathrm{x}_l) = \mathtt{true} \\ \{0\} & \text{otherwise} \end{cases}$$

For every $d_i^w \in V_2$ let $\mathcal{P}_2(d_i^w) \stackrel{def}{=} \{j_0 w\}$ where j_0 is the smallest index for which $\bigwedge_{k=1}^{K} \mathcal{A}(B_{i,j_0,k}) =$ true (such an index j_0 exists because $\mathcal{A}$ satisfies Φ). Obviously $\mathcal{P}_1, \mathcal{P}_2$ are non-empty, and the l_1 cost of the labeling is exactly $L + t \cdot W$.

Let us show that the labeling $(\mathcal{P}_1, \mathcal{P}_2)$ is a total cover. Let $e = (v_l, d_i^w)$ be an arbitrary edge, and let $jw \in \mathcal{P}_2(d_i^w)$. By definition of $\mathcal{P}_2$, j is such that $\mathcal{A}(B_{i,j,k}) =$ true for all $1 \leq k \leq K$. In addition, the edge e exists because x_l appears in the i^{th} disjunction (call these appearances $B_{i,j_1,k_1}, .., B_{i,j_R,k_R}$, $R > 0$). If $j \in \{j_1, .., j_R\}$ we use $(e, w, jw) \in \Pi$ to cover e, since $\mathcal{A}(B_{i,j=j_r,k_r}) =$ true and so $\mathcal{P}_1(v_l) = \{0, 1, .., W\}$ and in particular $w \in \mathcal{P}_1(v_l)$. If $j \notin \{j_1, .., j_R\}$ (i.e. x_l doesn't appear in the j^{th} conjunction of the i^{th} disjunction) we can cover e by $(e, 0, jw) \in \Pi$, since $0 \in \mathcal{P}_1(v_l)$ for every vertex $v_l \in V_1$. ■

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 $(\mathcal{P}_1, \mathcal{P}_2)$ be a total cover with l_1 cost $gt \cdot W$. Since $\forall v \in V_1 \quad \mathcal{P}_1(v) \subseteq \{0, 1, .., W\}$, and $\sum_{v \in V_1} |\mathcal{P}_1(v)| = gt \cdot W$, there must be at least one $w_0 > 0$ for which $|\{v \mid w_0 \in \mathcal{P}_1(v)\}| \leq gt$. We claim that the (weight-gt) assignment $\mathcal{A}$ defined by assigning x_l the value true if and only if $w_0 \in \mathcal{P}_1(v_l)$, satisfies Φ:

A label $b \in \mathcal{P}_2(d_i^{w_0}) \neq \phi$ must be an integer multiple of w_0, otherwise it cannot be covered. We will show that for every i, the j^{th} conjunction (where $j = b/w_0$) is satisfied. For this purpose we need to show that the literal x_{l_k} represented by $B_{i,j,k}$ is assigned the value true for every k, or $w_0 \in \mathcal{P}_1(v_{l_k})$. This is immediate since there is no other way of covering the edges $e_k \stackrel{def}{=} (v_{l_k}, d_i^{w_0})$, and $(\mathcal{P}_1, \mathcal{P}_2)$ is a total-cover.

■

Summing up 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 $> g \cdot tW$. Thus 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 directly.

4 Discussion and Open Questions

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 $i \geq 3$?

A Depend-2 $\mathcal{PCP}$ 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 after fixing an edge, the relation $\Pi \subset E \times B_1 \times B_2$ is a partial function. Our result inherently cannot be extended to this version, the main obstacle being that there is no known $\mathcal{PCP}$ characterization of NP with *exactly two* provers (i.e. a $\mathcal{PCP}$ test-system where each tests accesses exactly two variables, called depend-2-$\mathcal{PCP}$). Compare this to the known low error-probability $\mathcal{PCP}$ characterization of NP [RS97, DFK$^+$] 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.

For proving the hardness of approximating the CLOSEST-VECTOR problem, this obstacle was bypassed [DKS98] by showing NP-hardness for a different problem called $\mathcal{SSAT}$. $\mathcal{SSAT}$ could then be used instead of depend-2-$\mathcal{PCP}$ for proving NP-hardness for approximating the CLOSEST-VECTOR problem (and other problems) to within large factors. The main difference between LABEL-COVER and $\mathcal{SSAT}$ is that $\mathcal{SSAT}$ deals with *integer linear combinations* i.e. allowing 'negative-weights', while LABEL-COVER deals with *subsets* i.e. allowing only 'positive-weights'. $\mathcal{SSAT}$ and LABEL-COVER make up two ways of overcoming the lack of a depend-2 $\mathcal{PCP}$ characterization of NP.

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.

[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+] I. Dinur, E. Fischer, G. Kindler, R. Raz, and S. Safra. PCP characterizations of NP: Towards a polynomially-small error-probability. Manuscript, to appear in STOC99.

[DKRS99] I. Dinur, G. Kindler, R. Raz, and S. Safra. Approximating-CVP to within almost-polynomial factors is NP-hard. Manuscript, 1999.

[DKS98] I. Dinur, G. Kindler, and S. Safra. Approximating-CVP to within almost-polynomial factors is NP-hard. In *Proc. 39th IEEE Symp. on Foundations of Computer Science*, 1998.

[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.

[Hoc97] (Ed.) D. S. Hochbaum. Approximation algorithms for NP-hard problems (PWS publishing company, 1997) (contents). *SIGACTN: SIGACT News (ACM Special Interest Group on Automata and Computability Theory)*, 28, 1997.

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

[KST97] Sanjeev Khanna, Madhu Sudan, and Luca Trevisan. Constraint satisfaction: The approximability of minimization problems. In *Proceedings, Twelfth Annual IEEE Conference on Computational Complexity*, pages 282–296, Ulm, Germany, 24–27 June 1997. IEEE Computer Society Press.

[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.

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'