

# COMBINATION CAN BE HARD: APPROXIMABILITY OF THE UNIQUE COVERAGE PROBLEM\*

ERIK D. DEMAINE<sup>†</sup>, URIEL FEIGE<sup>‡</sup>, MOHAMMADTAGHI HAJIAGHAYI<sup>†</sup>, AND  
MOHAMMAD R. SALAVATIPOUR<sup>§</sup>

**Abstract.** We prove semilogarithmic inapproximability for a maximization problem called *unique coverage*: given a collection of sets, find a subcollection that maximizes the number of elements covered exactly once. Specifically, assuming that  $\text{NP} \not\subseteq \text{BPTIME}(2^{n^\varepsilon})$  for an arbitrary  $\varepsilon > 0$ , we prove  $O(1/\log^\sigma n)$  inapproximability for some constant  $\sigma = \sigma(\varepsilon)$ . We also prove  $O(1/\log^{1/3-\varepsilon} n)$  inapproximability for any  $\varepsilon > 0$ , assuming that refuting random instances of 3SAT is hard on average; and we prove  $O(1/\log n)$  inapproximability under a plausible hypothesis concerning the hardness of another problem, balanced bipartite independent set. We establish an  $\Omega(1/\log n)$ -approximation algorithm, even for a more general (budgeted) setting, and obtain an  $\Omega(1/\log B)$ -approximation algorithm when every set has at most  $B$  elements. We also show that our inapproximability results extend to envy-free pricing, an important problem in computational economics. We describe how the (budgeted) unique coverage problem, motivated by real-world applications, has close connections to other theoretical problems, including max cut, maximum coverage, and radio broadcasting.

**Key words.** unique coverage, hardness of approximation, wireless networks

**AMS subject classifications.** 68Q25, 68W25

**1. Introduction.** In this paper we consider the approximability of the following natural maximization analogue of the set cover problem:

**Unique Coverage Problem.** Given a universe  $U = \{e_1, \dots, e_n\}$  of elements, and given a collection  $\mathcal{S} = \{S_1, \dots, S_m\}$  of subsets of  $U$ , find a subcollection  $\mathcal{S}' \subseteq \mathcal{S}$  to maximize the number of elements that are *uniquely covered*, i.e., appear in exactly one set of  $\mathcal{S}'$ .

We also consider a generalized form of this problem that is useful for several applications (detailed in Section 2):

**Budgeted Unique Coverage Problem.** Given a universe  $U = \{e_1, \dots, e_n\}$  of elements, a profit  $p_i$  for each element  $e_i$ , a collection  $\mathcal{S} = \{S_1, \dots, S_m\}$  of subsets of  $U$ , a cost  $c_i$  of each subset  $S_i$ , and a budget  $B$ , find a subcollection  $\mathcal{S}' \subseteq \mathcal{S}$ , whose total cost is at most the budget  $B$ , to maximize the total profit of elements that are *uniquely covered*, i.e., appear in exactly one set of  $\mathcal{S}'$ .

*Motivation.* Logarithmic inapproximability for minimization problems is by now commonplace, starting in 1993 with a result for the celebrated set cover problem [42], which has since been improved to the optimal constant [18] and to assume just  $\text{P} \neq \text{NP}$  [47], and has been used to prove other tight (not necessarily logarithmic) inapproximability results for a variety of minimization problems; see, e.g., [33, 25, 13]. In contrast, for maximization problems,  $\log n$  inapproximability seems more difficult, and relatively few results are known. The only three such results of which we are aware are

---

\*A preliminary version of this paper appeared in *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2006.

<sup>†</sup>Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 (edemaine@mit.edu, hajiagha@mit.edu).

<sup>‡</sup>Department of Computer Science and Applied Mathematics, Weizmann Institute, Rehovot 76100, Israel (uriel.feige@weizmann.ac.il).

<sup>§</sup>Department of Computing Science, University of Alberta, Edmonton, AB, T6G 2E8, Canada (mreza@cs.ualberta.ca). Supported by NSERC and a faculty startup grant.

$(1 + \varepsilon)/\ln n$  inapproximability for domatic number unless  $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$  [21];  $1/\log^{1/2-\varepsilon} n$  inapproximability for the maximum edge-disjoint paths problems unless  $\text{NP} \subseteq \text{ZPTIME}(n^{\text{polylog } n})$  [5, 6, 48]; and  $1/\log^{1-\varepsilon} n$  inapproximability for the maximum edge-disjoint cycles problems unless  $\text{NP} \subseteq \text{ZPTIME}(n^{\text{polylog } n})$  [38]. Although these problems are interesting, they are rather specific, and we lack a central maximization problem analogous to set cover to serve as a basis for further reduction to many other maximization problems.

The unique coverage problem defined above is a natural maximization version of set cover which was brought to our attention from its applications in wireless networks. In one (simplified) application, we have a certain budget to build/place some transmitters at a subset of some specified set of possible locations. Our goal is to maximize the clients that are “covered” by (i.e., are within the range of) exactly one transmitter; these are the clients that receive signal without interference; see Section 2.1 for details. Another closely related application is the radio broadcast problem, in which a message (starting from one node of the network) is to be sent to all the nodes in the network in rounds. In each round, some of the nodes that have already received the message send it to their neighbors, and a node receives a message only if it receives it from exactly one of its neighbors. The goal is to find the minimum number of rounds to broadcast the message to all the nodes; see Section 2.5 for details. Therefore, every single round of a radio broadcast can be seen as a unique coverage problem. These applications and others are studied in more detail in Section 2.

*Known results.* To the best of our knowledge, there is no explicit study in the literature of the unique coverage problem and its budgeted version. However, the closely related radio broadcast problem has been studied extensively in the past, and implicitly includes an  $\Omega(1/\log n)$ -approximation algorithm for the basic (unbudgeted) unique coverage problem; see Section 2.5 for details.

Concurrently and independently of our work, Guruswami and Trevisan [28] study the so called 1-in- $k$  SAT problem, which includes the unique coverage problem (but not its budgeted version) as a special case. In particular, they show that there is an approximation algorithm that achieves an approximation ratio of  $1/e$  on satisfiable instances (in which all items can be covered by mutually disjoint sets).

*Our results.* On the positive side, we give an  $\Omega(1/\log n)$ -approximation for the budgeted unique coverage problem. We also show that, if each set has a bound  $B$  on the ratio between the maximum profit of a set (where the profit of a set is the sum of profits of its elements) and the minimum profit of an element, then budgeted unique coverage has an  $\Omega(1/\log B)$ -approximation. Section 4 proves these results.

The main focus of this paper is proving the following inapproximability results. We show that it is hard to approximate the unique coverage problem within a factor of  $\Omega(1/\log^\sigma n)$  for some constant  $\sigma$  depending on  $\varepsilon$ , assuming that  $\text{NP} \not\subseteq \text{BPTIME}(2^{n^\varepsilon})$  for some constant  $\varepsilon > 0$ . This inapproximability can be strengthened to  $\Omega(1/\log^{1/3-\varepsilon} n)$  (for any constant  $\varepsilon > 0$ ) under the assumption that refuting random instances of 3SAT is hard on average (hardness of R3SAT as in [19]). The inapproximability can be further strengthened to  $1/(\varepsilon \log n)$  for some constant  $\varepsilon > 0$ , under a plausible hardness hypothesis about a problem called Balanced Bipartite Independent Set; see Hypothesis 3.22. Section 3 proves all of these results.

Our hardness results have other implications regarding the hardness of some well-studied problems. In particular, for the problem of unlimited-supply single-minded (envy-free) pricing, a recent result [27] proves an  $\Omega(1/\log n)$ -approximation, but no inapproximability result better than APX-hardness is known. As we show in Section 2.2,

our hardness results for the unique coverage problem imply the same hardness-of-approximation bounds for this version of envy-free pricing. For the radio broadcast problem, as we discuss in Section 2.5, there is essentially a gap of  $\Omega(\log n)$  between the approximation and inapproximability factors ( $O(\log^2 n)$  vs.  $\Omega(\log n)$ ). We believe that our technique to prove hardness of unique coverage may shed some light on how to obtain a hardness of approximation beyond  $\Omega(\log n)$  for this problem.

More generally, there are many maximization problems for which the best known approximation factor is  $\Omega(1/\log n)$ —see, e.g., [27, 9, 43]—and it is not known whether an  $\Omega(1)$ -factor approximation is possible. Often (as indeed is the case with unique coverage) these problems naturally decompose into  $\Theta(\log n)$  subproblems, where at least an  $\Omega(1/\log n)$  fraction of the optimum’s value comes from one of these subproblems. In isolation, each subproblem can be approximated up to a constant factor, leading to an  $\Omega(1/\log n)$ -approximation algorithm for the whole problem. It may appear that this isolation approach is too naïve to give the best possible approximation, and that by a clever combination of the subproblems, it should be possible to get an  $\Omega(1)$ -approximation algorithm. Our hardness results show to the contrary that such intelligent combination can be hard, in the sense that the naïve isolation approach cannot be substantially improved, and suggest how one might obtain better hardness results for these problems.

## 2. Applications and Related Problems.

**2.1. Wireless Networks.** Our original motivation for the budgeted unique coverage problem is a real-world application arising in wireless networks.<sup>1</sup> We are given a map of the densities of mobile clients throughout a service region (e.g., the plane with obstacles). We are also given a collection of candidate locations for wireless base stations, each with a specified building cost and a specified coverage region (typically a cone or a disk, possibly obstructed by obstacles). This collection may include multiple options for base stations at the same location, e.g., different powers and different orientations of antennas. The goal is to choose a set of base stations and options to build, subject to a budget on the total building cost, in order to maximize the density of served clients.

The difficult aspect of this problem (and what distinguishes it from maximum coverage—see Section 2.4) is interference between base stations. In the simplest form, there is a limit  $k$  on the number of base stations that a mobile client can reasonably hear without conflict between the signals; any client within range of more than  $k$  base stations cannot communicate because of interference and thus is not serviced. More generally, a mobile client’s reception is better when it is within range of fewer base stations, and our goal is to maximize total reception. To capture these desires, the instance specifies the *satisfaction*  $s_i$  of a client within range of exactly  $i$  base stations, such that  $s_0 = 0$  and  $s_1 \geq s_2 \geq s_3 \geq \dots \geq 0$ . The goal is to choose a set of base stations and options, again subject to the budget constraint, in order to maximize the total satisfaction weighted by client densities.

When all  $s_i$ ’s are equal, we just have the maximum coverage problem (Section 2.4). When  $s_1 = 1$  and  $s_i = 0$  for all  $i \neq 1$ , this problem can be formulated as a budgeted unique coverage problem, by standard discretization of the density map. More gen-

---

<sup>1</sup>The application arises in the context of cellular networks at Bell Labs. The problem we consider here is a somewhat simplified theoretical formulation of this application. In the real application, the interference patterns are more complicated, but this problem seems to be the cleanest theoretical formulation.

erally, for any assignment of  $s_i$ 's, the problem can be formulated as a generalization of budgeted unique coverage, the *budgeted low-coverage problem*. In this problem, we are also given satisfaction factors  $s_i$  for an element being covered exactly  $i$  times, zero for  $i = 0$  and non-increasing for  $i > 0$ , and the goal is to maximize the total satisfaction, i.e., the sum over all elements of the product of the element's profit (here, density) and its satisfaction factor (the appropriate  $s_i$ ). We show that our approximation algorithms for the budgeted unique coverage problem apply more generally to the budgeted low-coverage problem, yielding an  $\Omega(1/\log n)$ -approximation where  $n$  is the total number of options for base stations.

While similar problems about base-station placement have been considered before, very few works consider maximization forms of the problem, which is the focus of this paper. Lev-Tov and Peleg [41] consider the following very specialized form of the problem: base stations are unit disks in the plane, and the goal is to maximize the number of uniquely receiving clients. For this problem they give an  $n^{O(\sqrt{n})}$ -time algorithm, where  $n$  is the number of candidate disks. In the application of interest, we believe that it is more natural to allow clients to be covered more than once but reduce (or eliminate) the satisfaction of these clients; this removal of an artificial constraint may enable substantially better solutions to the problem. Other work [31, 22, 8] solves the problem of assigning powers to base stations such that, when each client prefers a unique base station, we do not violate the capacities of the base stations, provided the number of clients is at most the total capacity of the network.

**2.2. Envy-Free Pricing.** Fundamental to “fair” equilibrium pricing in economics is the notion of envy-free pricing [49, 26]. This concept has recently received attention in computer science [1, 27] in the new trend toward an algorithmic understanding of economic game theory; see, e.g., [14, 15] for related work.

The following version of envy-free pricing was considered in [27]. A single seller prices  $m$  different items,  $I_1, \dots, I_m$ , each with a specified quantity (limited or unlimited *supply*). Each of  $n$  buyers  $b_i$  ( $1 \leq i \leq n$ ) wishes to purchase a subset of items (a *bundle*), and the seller knows the maximum price that each buyer is willing to pay for each bundle (the *valuation*). A buyer's *utility* is the difference between the valuation and the price of the bundle (sum of the prices of the items in the bundle) as sold to the buyer. The seller must choose the item prices, price  $p_i$  for item  $I_i$ , and which bundles are sold to which buyers in such a way that is *envy-free*: each buyer should be sold a bundle that has the maximum utility among all bundles. The goal is to maximize the seller's *profit*, i.e., the total price of the sold bundles.

Among other results, Guruswami et al. [27] give an  $\Omega(1/(\log n + \log m))$ -approximation algorithm for the unlimited-supply *single-minded* bidder problem, where each buyer  $b_i$  considers only one particular bundle  $B_i$  and buys it if the cost is less than the valuation. They also give a constant-factor hardness-of-approximation result for this problem, via a reduction from max-cut. Single-minded bidders were considered before in the context of combinatorial auctions and mechanism design [7, 45, 40]. The unlimited-supply assumption in combination with single-mindedness simplifies the problem, as the notion of *envy* does not play a role in this case. The general version of the envy-free pricing problem is of course at least as difficult as this special case.

We now show that unlimited-supply single-minded (envy-free) pricing is as hard to approximate as the unique coverage problem. The reduction is as follows. Each set  $S_i$  in the collection (for the instance of unique coverage) maps to an item  $I_i$  (for the instance of envy-free pricing). Each element  $e_i$  of the universe  $U$  maps to a buyer  $b_i$ . Buyer  $b_i$  has a valuation of 1 for one bundle,  $B_i$ , namely, the set of

items  $I_j$  that correspond to sets  $S_j$  containing the element  $e_i$ . In this context, every price assignment is envy-free, because we have unlimited supply for each item so the seller can always sell each buyer its desired bundle (if the buyer wants). Because the valuations are all 1, we can assume that all prices are between 0 and 1. By randomized rounding (see Lemma A.1), we can assume that all prices are either 0 or 1, at a loss of a constant factor in profit. In this case, each buyer  $b_i$  will buy its bundle precisely if at most one item is priced at 1, and the rest of the items are priced at 0. If all items in a bundle are priced at 0, then the seller makes no profit; if exactly one item is priced at 1 and the rest are priced at 0, then the seller profits by 1. Thus the effective goal is to assign prices of 0 or 1 in order to maximize the number of bundles for which exactly one item is priced at 1, which is identical to the original unique coverage problem.

Therefore our hardness-of-approximation results apply to unlimited-supply single-minded (envy-free) pricing and establish semilogarithmic inapproximability.

**2.3. Max-Cut.** Recall the max-cut problem: given a graph  $G$ , find a cut  $(S, \bar{S})$ , where  $S \subseteq V(G)$  and  $\bar{S} = V(G) - S$ , that maximizes the number of edges with one endpoint in  $S$  and the other endpoint in  $\bar{S}$ . The max-cut problem can be seen to be equivalent to a special case of the unique coverage problem in which every element is in exactly two sets. Simply view every vertex as a set and every edge as an element.

Max-cut is 0.878567-approximable [24], 0.941177-inapproximable assuming  $P \neq NP$  [29], and 0.878568-inapproximable assuming the Unique Games Conjecture [34]. From these results one can immediately obtain constant-factor hardness for unique coverage, but in this paper we show that unique coverage is in fact much harder.

**2.4. Maximum Coverage.** Our budgeted unique coverage problem is also closely related to the *budgeted maximum coverage* variation of set cover: given a collection of subsets  $S$  of a universe  $U$ , where each element in  $U$  has a specified weight and each subset has a specified cost, and given a budget  $B$ , find a subcollection  $S' \subseteq S$  of sets, whose total cost is at most  $B$ , in order to maximize the total weight of elements covered by  $S'$ . For this problem, there is a  $(1 - 1/e)$ -approximation [30, 33], and this is the best constant approximation ratio possible unless  $P = NP$  [18, 33].

At first glance, one might expect the greedy  $(1 - 1/e)$ -approximation algorithm to work for unique coverage as well: the only difference between the two problems is whether we count elements that are covered (contained in at least one set) or uniquely covered (contained in exactly one set). Of course, we show that the (in)approximability of the two problems is quite different. Indeed, a natural class of greedy algorithms can be very bad for unique coverage. Consider the collection of sets  $S_i = \{i, k + 1, k + 2, \dots, n\}$  for  $i = 1, 2, \dots, k$ , with an infinite budget  $B$ . Consider a greedy algorithm that repeatedly chooses a set to add to the cover, according to some (arbitrary) rule, with one of two stopping conditions: either when the budget is exhausted, or when the number of uniquely covered elements goes down. Then the approximation ratio is  $\Theta(1/n)$  with the first stopping condition if  $k = 2$ , and with the second stopping condition if  $k = n - 2$ .

**2.5. Radio Broadcast.** The unique coverage problem is closely related to a single “round” of the *radio broadcast* problem [10]. This problem considers a *radio network*, i.e., a network of processors (nodes) that communicate synchronously in rounds. In each round, a node can either transmit to all of its neighbors in an undirected graph (representing the communicability between pairs of nodes) or not transmit. A node receives a message if exactly one of its neighbors transmits a message in the round; otherwise the messages are lost because of radio interference. In the

radio broadcast problem, initially one node has a message, and the goal is to propagate this message to all nodes in the network.

Radio broadcast is one of the most important communication primitives in radio networks, and the problem has been studied extensively in the literature. In summary, the current best algorithms for approximating the minimum number of rounds are a (multiplicative)  $O(\log^2 n)$ -approximation [12, 10, 36, 35] and an additive  $O(\log^2 n)$ -approximation [23]. Alon et al. [3] show that, even for graphs with diameter 3,  $\Omega(\log^2 n)$  rounds can be necessary. The problem has also been considered in the context of distributed algorithms [39, 37] and low-energy ad-hoc wireless networks [4]. Elkin and Kortsarz prove a lower bound of inapproximability of a (multiplicative)  $\Omega(\log n)$  [16] and an additive  $\Omega(\log^2 n)$  [17] assuming  $\text{NP} \not\subseteq \text{BPTIME}(n^{O(\log \log n)})$ .

The unique coverage problem (but not the budgeted version) can be considered as a single round of a greedy algorithm for the radio broadcast problem, which maximizes the number of nodes that receive the message in each step. Specifically, consider the bipartite subgraph where one side consists of all nodes that currently have the message and the other side consists of all nodes that do not yet have the message. In one round of the greedy algorithm, the goal is to find a subset of nodes in the first side to transmit in order to maximize the number of nodes in the second side that (uniquely) receive the message. This problem is equivalent to unique coverage, viewing nodes on the first side as sets and the nodes on the second side as elements of the universe.

One implication of the radio broadcasting work on unique coverage is an implicit  $\Omega(1/\log n)$ -approximation algorithm for the (unbudgeted) unique coverage problem. Namely, there is a randomized broadcasting algorithm that, in each round, guarantees transmission to an  $\Omega(1/\log r)$  fraction of the  $r$  neighbors of nodes that currently have the message. Because  $r$  is an obvious upper bound on the number of successful transmissions of the message, this result is an  $\Omega(1/\log r) = \Omega(1/\log n)$  approximation in this special case. See, e.g., [10].

To avoid the possibility of misunderstanding, let us point out that the known hardness-of-approximation results for radio broadcast [16, 17] do not give (neither explicitly nor implicitly) any useful hardness-of-approximation result for the unique coverage problem (not even a constant factor). Likewise, our hardness-of-approximation results for the unique coverage problem do not by themselves imply any new hardness-of-approximation results for radio broadcast. However, they do introduce a component that may be useful in future hardness-of-approximation results for the radio broadcast problem, as they show that the greedy broadcast policy might need to lose a semilogarithmic factor already in a single round (a fact not used in [16, 17]).

**3. Inapproximability.** In this section we prove that it is hard to approximate unique coverage within a factor of  $\Omega(1/\log^c n)$  for some constant  $c$ ,  $0 < c \leq 1$ . Our main result is a general reduction from a variation of the Balanced Bipartite Independent Set (BBIS) problem (defined below) to the unique coverage problem. From this reduction and the known hardness results for BBIS, we can derive an  $O(1/\log^c n)$  hardness for unique coverage. Under a plausible assumption about the hardness of BBIS, this bound can be improved to  $O(1/\log n)$ .

We consider the natural graph-theoretic model of the unique coverage problem. Define the bipartite graph  $H(V \cup W, F)$  with a vertex  $v_i \in V$  for every set  $S_i \in \mathcal{S}$  and a vertex  $w_j \in W$  for every element  $e_j \in U$ , and with an edge  $f = (v_i, w_j) \in F$  precisely if  $e_j \in S_i$ . Then the unique coverage problem asks to find a subset  $V' \subseteq V$  such that the subgraph induced by  $V' \cup W$  has the maximum number of degree-1

vertices in  $W$ . We call the degree-1 vertices *uniquely covered* by the vertices in  $V'$ .

**DEFINITION 3.1.** *Given a bipartite graph  $G(A \cup B, E)$  with  $|A| = |B| = n$ , the *Balanced Bipartite Independent Set (BBIS)* problem asks to find the largest value of  $k$  such that there are sets  $\tilde{A} \subseteq A$  and  $\tilde{B} \subseteq B$  with  $|\tilde{A}| = |\tilde{B}| = k$  where the subgraph  $\tilde{G}$  of  $G$  induced by  $\tilde{A} \cup \tilde{B}$  is an independent set.*

As detailed below, this problem has known hardness results (see [19, 32]). In order to prove hardness of the unique coverage problem, we define a variation of BBIS. Then we give a reduction from this variation of BBIS. Before stating the main result, we need to define what we mean by an  $(a, b)$ -BIS (bipartite independent set). Let  $G(A \cup B, E)$  be a given a bipartite graph. If the subgraph  $\tilde{G}$  induced by  $\tilde{A} \subseteq A$  and  $\tilde{B} \subseteq B$ , with  $|\tilde{A}| = a$  and  $|\tilde{B}| = b$ , is an independent set, then we call it an  $(a, b)$ -BIS.

**DEFINITION 3.2.** *Given bipartite graph  $G(A \cup B, E)$  with  $|A| = |B| = n$ , and given parameters  $\gamma, \gamma', \delta$ , and  $\delta'$  satisfying  $0 < \gamma' < \gamma \leq 1$  and  $0 \leq \delta < \delta' \leq 1$ , the  $\text{BBIS}(\gamma, \gamma', \delta, \delta')$  problem is to distinguish between two cases:*

1. Yes instance:  $G$  has an  $(n^\gamma, n/\log^\delta n)$ -BIS.
2. No instance:  $G$  has no  $(n^{\gamma'}, n/\log^{\delta'} n)$ -BIS.

The main theorem of this section is the following:

**THEOREM 3.3.** *There is a polynomial-time probabilistic reduction from BBIS to the unique coverage problem with the following properties. Given a bipartite graph  $G(A \cup B, E)$  with  $|A| = |B| = n$  and given constants  $\gamma, \gamma', \delta$ , and  $\delta'$  satisfying  $0 < \gamma' < \gamma \leq 1$  and  $0 \leq \delta < \delta' \leq 1$ , the algorithm constructs in randomized polynomial time an instance  $H(V \cup W, F)$  of unique coverage with  $|W| = \Theta((\gamma - \gamma')n \log n)$  and  $|V| = n$  satisfying the following two properties:*

1. If  $G$  is a Yes instance of  $\text{BBIS}(\gamma, \gamma', \delta, \delta')$ , then  $H$  has a solution of size  $\Omega((\gamma - \gamma')n \log^{1-\delta} n)$ .
2. If  $G$  is a No instance of  $\text{BBIS}(\gamma, \gamma', \delta, \delta')$ , then  $H$  has no solution of size  $O((\gamma - \gamma')n \log^{1-\delta'} n)$ .

**COROLLARY 3.4.** *Assuming that  $\text{BBIS}(\gamma, \gamma', \delta, \delta')$  is hard for constants  $\gamma, \gamma', \delta, \delta'$ , we get a hardness of approximation within a factor of  $\Omega(1/\log^{\delta'-\delta} n)$  for unique coverage.*

This theorem is proved in Section 3.1. Next we show how the known hardness results for BBIS can be used to derive explicit hardness results for unique coverage. In particular, the following theorems follow from Theorem 3.3.

**THEOREM 3.5.** *Let  $\varepsilon > 0$  be an arbitrarily small constant. Assuming that  $\text{NP} \not\subseteq \text{BPTIME}(2^{n^\varepsilon})$ , it is hard to approximate the unique coverage problem within a factor of  $\Omega(1/\log^\sigma n)$  for some constant  $\sigma = \sigma(\varepsilon)$ .*

Feige [19] makes the following hypothesis about average-case hardness of 3SAT.

**R3SAT hardness hypothesis [19]:** Let  $\phi$  be a 3SAT formula with  $n$  variables and  $m = \Delta n$  clauses where every clause is generated independently at random by selecting three literals independently at random. For arbitrary large constant  $\Delta$ , there is no polynomial time algorithm that accepts if  $\phi$  is satisfiable and refutes at least half of the times for those formulas that are not satisfiable.

Under this hypothesis, we can prove the same hardness result with an explicit value for  $\sigma$ :

**THEOREM 3.6.** *Assuming R3SAT hardness hypothesis, unique coverage is hard to approximate within a factor of  $\Omega(1/\log^{1/3-\sigma} n)$  for an arbitrarily small constant  $\sigma > 0$ .*

Under a stronger (yet plausible) hardness assumption, explained in Section 3.2, we close the gap between the approximation factor and the hardness of approximation, up to the constant multiplicative factor, by proving an  $O(1/\log n)$ -hardness result for unique coverage.

**THEOREM 3.7.** *Assuming a specific hardness of factor  $\Omega(n^\varepsilon)$  for BBIS for some constant  $\varepsilon > 0$  (Hypothesis 3.22), it is hard to approximate the unique coverage problem within a factor of  $\Omega(1/\log n)$  where the constant in the  $\Omega$  term depends on  $\varepsilon$ .*

Theorems 3.5 to 3.7 are proved in Section 3.2.

### 3.1. Reduction from BBIS to Unique Coverage and Proof of Theorem 3.3.

**Construction.** Consider an instance of BBIS( $\gamma, \gamma', \delta, \delta'$ ): a bipartite graph  $G(A \cup B, E)$  with  $|A| = |B| = n$ , and constants  $\gamma, \gamma', \delta$ , and  $\delta'$  with  $0 < \gamma' < \gamma \leq 1$  and  $0 \leq \delta < \delta' \leq 1$ . We construct a graph  $H(V \cup W, F)$  as an instance of unique coverage as follows.

First we construct a random graph  $G'(A' \cup B', E')$  where  $A'$  is a copy of  $A$  and  $B'$  is a copy of  $B$ . For every  $a \in A'$  and  $b \in B'$  we place the edge  $(a, b)$  in  $E'$  with probability  $1/n^\gamma$ . So the expected degree of every vertex in  $G'$  is  $n^{1-\gamma}$ .

Now to construct  $H$ , let  $V$  be a copy of  $A$ . Then with  $\gamma'' = \frac{\gamma-\gamma'}{\gamma}$ , create  $p = \gamma'' \log n$  copies of  $B$ , named  $W_1, \dots, W_p$ . We define a bipartite graph  $H_i(V \cup W_i, F_i)$ , for every  $1 \leq i \leq p$ , and at the end  $H = \bigcup_{i=1}^p H_i$ . Note that  $|V| = n$  and  $|W| = pn$ . The set of edges  $F_i$  (in  $H_i$ ) consists of the union of two edge sets: (i) the edges of the random graph  $G'$  induced on the vertices  $V \cup W_i$  ( $V$  as  $A'$  and  $W_i$  as  $B'$ ), plus (ii) the edges of another random graph  $G_i$ , where  $G_i$  is defined recursively as follows. Initially,  $G_1$  is  $G$  induced on  $V \cup W_1$ . For every  $i \geq 2$ ,  $G_i$  is obtained from  $G_{i-1}$  by deleting every edge independently with probability  $\frac{1}{2}$ . The edges of  $G'$  in  $H_i$  are called *type-1 edges* and the rest of the edges of  $H_i$  (which come from  $G_i$ ) are called *type-2 edges* of  $H_i$ . In a solution to  $V' \subseteq V$  for instance  $H$ , we say a vertex of  $W$  is uniquely covered by a type-1 (type-2) edge if that vertex is adjacent to exactly one vertex of  $V'$  and that edge is a type-1 (type-2) edge.

**Proof overview.** Here is the general idea of the proof. Intuitively, the balanced independent sets in  $G$  relate to the elements that will be uniquely covered by type-1 edges in  $H$ . The removal of edges (randomly) from  $G_i$  to  $G_{i+1}$  is to ensure that not too many vertices are uniquely covered by type-2 edges. More specifically, we will show that the number of vertices uniquely covered by type-2 edges (edges that were originally in  $G$ ) in this instance is  $O(n)$ . So let us focus on the vertices uniquely covered by type-1 edges (i.e., edges from the random graph  $G'$  in each  $H_i$ ).

First suppose that  $G$  is a Yes instance, i.e., it has an  $(n^\gamma, \frac{n}{\log^\delta n})$ -BIS, say  $A^* \cup B^*$  (with  $A^* \subseteq A$  and  $B^* \subseteq B$ ). Because the expected degree of every vertex in  $G'$  is  $n^{1-\gamma}$ , the expected number of type-1 edges coming out of  $A^*$  (in  $G'$ ) is  $n$ , and because these edges are selected at random, we expect a fraction of  $1/e$  of the vertices in  $B'$  (in  $G'$ ) and in particular a fraction of  $1/e$  of the vertices in  $B^*$  to have degree 1. This implies that the type-1 edges in each  $H_i$  uniquely cover a linear number of vertices of  $B^*$  (at least in expectation), i.e., it gives a solution of size  $\Omega(\frac{n}{\log^\delta n})$  in  $H_i$ . Because  $H = \bigcup_{i=1}^p H_i$  and  $p = \gamma'' \log n$ , we have a total of  $\Omega(\gamma'' n \log^{1-\delta} n)$  vertices uniquely covered by type-1 edges.

Now suppose that  $G$  is a No instance, i.e., it has no  $(n^{\gamma'}, n/\log^{\delta'} n)$ -BIS. We will show that, although we delete edges to construct  $G_i$  from  $G_{i-1}$ , the last (and



most sparse) graph  $G_p$  will not have “too large” a bipartite independent set with high probability. This property will be used to show that, in every graph  $H_i$ , the number of vertices uniquely covered by type-1 edges in any solution of  $H$  is at most  $O(n/\log^{\delta'} n)$  with high probability. Thus, the total number of vertices uniquely covered in  $H$  (by type-1 or type-2 edges) in any solution is at most  $O(\gamma'' n \log^{1-\delta'} n + n)$  with high probability. Because  $\delta' \leq 1$ , this creates a hardness gap of  $\Omega(1/\log^{\delta'-\delta} n)$ .

Now we give the details of the proof of Theorem 3.3. We use the following simplified version of the Chernoff bound:

LEMMA 3.8 (Chernoff bound). *For independent 0/1 random variables  $X_1, \dots, X_n$ ,  $X = \sum_{i=1}^n X_i$ ,  $\mu = E[X]$ , and any  $0 < \delta < 1$ , we have*

$$\Pr[|X - E[X]| > \delta\mu] \leq e^{-\delta^2\mu/3}.$$

LEMMA 3.9. *For every selection of vertices  $V' \subseteq V$  as a solution for instance  $H$ , the number of vertices uniquely covered by type-2 edges in any solution to  $H$  is  $O(n)$  with high probability.*

*Proof.* Let  $b \in B$  be an arbitrary vertex (in  $G$ ) and assume that  $w_1, \dots, w_p$  are its corresponding vertices in  $W_1, \dots, W_p$ . Consider any subset  $V' \subseteq V$ . Assuming that  $V'$  is a solution to unique coverage, we compute the probability that exactly  $i$  vertices out of  $w_1, \dots, w_p$  are uniquely covered by type-2 edges (of the vertices of  $V'$ ). Assume that  $j$  is the first index for which  $w_j$  is uniquely covered by a type-2 edge, and  $w_j, \dots, w_{j+i-1}$  are the copies that are uniquely covered by a type-2 edge. Because every edge is deleted with probability  $\frac{1}{2}$  from  $G_t$  to  $G_{t+1}$  (for  $1 \leq t < p$ ), the probability that a single edge survives  $i$  rounds is  $2^{-i}$ . Let  $X_b$  be the number of copies of  $b$  (from  $w_1, \dots, w_p$ ) that are uniquely covered by a type-2 edge (by the vertices of  $V'$ ) and define  $X = \sum_{b \in B} X_b$ . Therefore,

$$E[X] = \sum_{b \in B} E[X_b] = n \sum_{i=1}^p \frac{i}{2^i} \leq 3n.$$

Using the Chernoff bound (Lemma 3.8), we obtain

$$\Pr[X \geq 6n] \leq e^{-4n}.$$

Because there are  $2^n$  subsets  $V'$ , a union bound shows that the probability that, for at least one of those sets, the number of vertices in  $W$  that are uniquely covered by type-2 edges is  $\geq 6n$  is at most  $2^n \cdot e^{-4n} \leq e^{-\Omega(n)}$ . This completes the proof of the lemma.  $\square$

**Completeness.** Suppose that  $G$  is a Yes instance, i.e., it has a  $(n^\gamma, n/\log^\delta n)$ -BIS, say,  $A^* \cup B^*$  where  $A^* \subseteq A$  and  $B^* \subseteq B$ . Assume that  $V'$  and  $W'_i$  are the subsets of vertices in  $H_i$  and  $A''$  and  $B''$  are the subsets of vertices in  $G'$  corresponding to  $A^*$  and  $B^*$ , respectively. Because  $G_i$  is obtained from  $G$  by deleting edges, there are no type-2 edges in  $V' \cup W'_i$  in  $H_i$  (for any  $1 \leq i \leq p$ ). Therefore, every vertex  $w \in W'_i$  (for all values of  $1 \leq i \leq p$ ) has degree 1 if and only if the corresponding vertex  $w \in B''$  (in  $G'$ ) has degree 1. For every  $w \in B''$ , let  $X_w$  be a 0/1 random variable that is 1 if and only if  $w \in B''$  has degree 1 (and so  $w$  is uniquely covered by a type-1 edge in  $H_i$  for all  $1 \leq i \leq p$ ). With  $X = \sum_{w \in B''} X_w$ ,

$$E[X] = \sum_{w \in B''} \Pr[X_w = 1]$$

$$\begin{aligned}
&= |B''| \cdot \binom{|A''|}{1} \cdot \frac{1}{n^\gamma} \left(1 - \frac{1}{n^\gamma}\right)^{|A''|-1} \\
&\geq \frac{|B''|}{e} \geq \frac{n}{e \log^\delta n}.
\end{aligned}$$

A simple application of the Chernoff bound shows that  $\Pr[X \leq \frac{n}{6 \log^\delta n}] \leq e^{-\Omega(n/\log^\delta n)}$ . Therefore, if we select the subset of vertices in  $V$  (in  $H$ ) corresponding to  $A^*$  (in  $G$ ) then, with high probability, there are at least  $p \cdot \frac{n}{6 \log^\delta n} = \Omega(\gamma'' n \log^{1-\delta} n)$  vertices in  $W$  uniquely covered (by type-1 edges). Thus, we have proved the following:

**COROLLARY 3.10.** *If  $G$  is a Yes instance then, with high probability,  $H$  has a unique cover of size  $\Omega(\gamma'' n \log^{1-\delta} n)$ .*

**Soundness.** Suppose that  $G$  is a No instance, i.e., it has no  $(n^{\gamma'}, n/\log^{\delta'} n)$ -BIS. Our goal is to show that, with high probability, every solution to unique coverage for  $H$  has size  $O(\gamma'' n \log^{1-\delta'} n)$ . Because by Lemma 3.9 the number of vertices uniquely covered by type-2 edges is  $O(n)$ , we only need to prove that, with high probability, the number of vertices uniquely covered by type-1 edges is at most  $O(\gamma'' n \log^{1-\delta'} n)$ .

Consider any solution to unique coverage for  $H$ . By construction of the  $H_i$ 's, it is easy to see that, for every vertex  $b \in B$  (in  $G$ ), if the corresponding vertex in  $W_i$  is uniquely covered by a type-1 edge in  $H_i$ , then all the corresponding vertices of  $b$  in the  $W_j$ 's, for  $i \leq j \leq p$ , are also uniquely covered by a type-1 edge. Therefore, if we prove that the number of vertices uniquely covered by type-1 edges in  $H_p$  is upper bounded (with high probability) by  $O(n/\log^{\delta'} n)$ , then because  $p = \gamma'' \log n$ , we obtain the claimed upper bound for the total number of vertices uniquely covered by type-1 edges.

Suppose that  $V' \subseteq V$  and  $W' \subseteq W_p$  are such that all the vertices in  $W'$  are uniquely covered by  $V'$ , and the edges that cover them are all type-1 edges. It is easy to see that  $V' \cup W'$  must be a bipartite independent set in  $G_p$  (otherwise there is some type-2 edge incident to some vertex  $w \in W'$  and therefore  $w$  is not uniquely covered).

**LEMMA 3.11.** *If  $V' \cup W'$  (with  $V' \subseteq V$  and  $W' \subseteq W_p$ ) is a bipartite independent set in  $G_p$ , then with high probability, either  $|V'| < n^{(\gamma+\gamma')/2}$  or  $|W'| < 2n/\log^{\delta'} n$ , i.e.,  $G_p$  has no  $(n^{(\gamma+\gamma')/2}, 2n/\log^{\delta'} n)$ -BIS.*

*Proof.* Suppose that  $V' \subseteq V$  and  $W' \subseteq W_p$  satisfy  $|V'| = n^{(\gamma+\gamma')/2}$  and  $|W'| = 2n/\log^{\delta'} n$ . Partition  $V'$  into  $q = n^{(\gamma-\gamma')/2}$  subsets  $V'_1, \dots, V'_q$ , each of size  $n^{\gamma'}$ . Let  $A_i^*$  and  $B^*$  ( $1 \leq i \leq q$ ) be the subset of vertices of  $A$  and  $B$  (in  $G$ ) corresponding to  $V'_i$  and  $W'$ , respectively. Consider the subgraph of  $G$  induced by  $A_i^* \cup B^*$ . Because  $|A_i^*| = n^{\gamma'}$ ,  $|B^*| = 2n/\log^{\delta'} n$ , and because  $G$  has no  $(n^{\gamma'}, n/\log^{\delta'} n)$ -BIS, it follows that at least  $n/\log^{\delta'} n$  vertices in  $B^*$  must be connected to the vertices in  $A_i^*$ . Therefore, the total number of edges in the subgraph induced by  $B^* \cup \bigcup_{i=1}^q A_i^*$  is at least  $q \cdot n/\log^{\delta'} n = \Omega(n^{1+(\gamma-\gamma')/2}/\log^{\delta'} n)$ . Because  $G_1 = G$ ,  $V' \cup W'$  forms an independent set in  $G_p$  only if all of these  $\Omega(n^{1+(\gamma-\gamma')/2}/\log^{\delta'} n)$  edges are deleted while  $G_p$  is created. Because in creating  $G_{i+1}$  from  $G_i$ , edges are deleted with probability  $\frac{1}{2}$ , we have

$$(3.1) \quad \Pr[V' \cup W' \text{ is an independent set in } G_p] \leq (1 - 2^{-p})^{\Omega(n^{1+(\gamma-\gamma')/2}/\log^{\delta'} n)}.$$

The number of such subsets  $V' \cup W'$  is

$$(3.2) \quad \binom{n}{n^{(\gamma+\gamma')/2}} \binom{n}{2n/\log^{\delta'} n}.$$

Thus, using (3.1) and (3.2), the expected number of bipartite independent sets  $V' \cup W'$  with  $|V'| = n^{(\gamma+\gamma')/2}$  and  $|W'| = 2n/\log^{\delta'} n$  in  $G_p$  is at most

$$\begin{aligned} & (1 - 2^{-p})^{\Omega(n^{1+(\gamma-\gamma')/2}/\log^{\delta'} n)} \binom{n}{n^{(\gamma+\gamma')/2}} \binom{n}{2n/\log^{\delta'} n} \\ & \leq (1 - n^{-(\gamma-\gamma')/7})^{\Omega(n^{1+(\gamma-\gamma')/2}/\log^{\delta'} n)} \cdot \left(\frac{en}{n^{(\gamma+\gamma')/2}}\right)^{n^{(\gamma+\gamma')/2}} \cdot \left(\frac{en}{2n/\log^{\delta'} n}\right)^{2n/\log^{\delta'} n} \\ & \leq e^{-\Omega(n^{1+(\gamma-\gamma')/2-(\gamma-\gamma')/7}/\log^{\delta'} n)} \cdot e^{O(n^{(\gamma+\gamma')/2} \log n)} \cdot e^{O(n \log \log n / \log^{\delta'} n)} \\ & \leq e^{-\Omega(n^{1+(\gamma-\gamma')/3})} \cdot e^{O(n/\log^{\delta'/2} n)} \\ & \leq e^{-\Omega(n^{1+(\gamma-\gamma')/3})}. \end{aligned}$$

Therefore, with probability  $1 - e^{-\Omega(n^{1+(\gamma-\gamma')/3})}$ , for every bipartite independent set  $V' \cup W'$  of  $G_p$ , either  $|V'| < n^{(\gamma+\gamma')/2}$  or  $|W'| < 2n/\log^{\delta'} n$ , i.e.,  $G_p$  has no  $(n^{(\gamma+\gamma')/2}, 2n/\log^{\delta'} n)$ -BIS.  $\square$

**LEMMA 3.12.** *With high probability, for every selection of vertices  $V' \subseteq V$  as a solution to instance  $H$ , for every  $H_i$  ( $1 \leq i \leq p$ ), the number of vertices uniquely covered by type-1 edges is at most  $O(n/\log^{\delta'} n)$ .*

*Proof.* Let  $V' \subseteq V$  be any fixed solution to instance  $H$ . Clearly, for every vertex uniquely covered by a type-1 edge in  $W_i$ , its corresponding copy is also uniquely covered (by a type-1 edge) in  $W_j$  for every  $i \leq j \leq p$ . So let us focus on the number of vertices uniquely covered by type-1 edges in  $H_p$ . If  $W'_p \subseteq W_p$  is the set of vertices uniquely covered by type-1 edges in  $H_p$  (by  $V'$ ), then  $V' \cup W'_p$  is a  $(|V'|, |W'_p|)$ -BIS in  $G_p$ . We are going to use Lemma 3.11 to prove that, with high probability (for all possible choices of  $V'$ ), the size of  $W'_p$  is small.

First consider the case that  $|V'| \geq n^{(\gamma+\gamma')/2}$  and  $V', W'_p$  form a  $(|V'|, |W'_p|)$ -BIS in  $G_p$ . By Lemma 3.11, the probability that  $|V'| \geq n^{(\gamma+\gamma')/2}$  and  $|W'_p| \geq 2n/\log^{\delta'} n$  is at most  $e^{-\Omega(n^{1+(\gamma-\gamma')/3})}$ . The number of solutions to  $H$  (i.e. subsets  $V'$ ) that satisfy the bound on  $V'$  is clearly at most  $2^n$ . Thus, the probability that there is a solution  $V'$  such that  $|V'| \geq n^{(\gamma+\gamma')/2}$  and  $|W'_p| \geq 2n/\log^{\delta'} n$  and  $V', W'_p$  forms a  $(|V'|, |W'_p|)$ -BIS in  $G_p$  is at most

$$(3.3) \quad 2^n \cdot e^{-\Omega(n^{1+(\gamma-\gamma')/3})} = e^{-\Omega(n^{1+(\gamma-\gamma')/3})}.$$

Now consider the case that  $|V'| < n^{(\gamma+\gamma')/2}$  (and of course  $|W'_p| \leq n$ ). In this case, we show that, with high probability,  $|W'_p| \leq O(n^{1-(\gamma-\gamma')/3})$ , which is clearly  $O(n/\log^{\delta'} n)$ . Consider an arbitrary vertex  $w \in W'_p$  and let  $X_w$  be a 0/1 random variable that is 1 if and only if  $w$  is incident to exactly one type-1 edge. With  $X = \sum_{w \in W'_p} X_w$ ,

$$E[X] = \sum_{w \in W'_p} \Pr[X_w = 1]$$

$$\begin{aligned}
&= \sum_{w \in W'_p} \binom{|V'|}{1} \cdot \frac{1}{n^\gamma} \left(1 - \frac{1}{n^\gamma}\right)^{|V'|-1} \\
&< n \cdot n^{(\gamma+\gamma')/2} \cdot \frac{1}{n^\gamma} \\
&\leq n^{1-(\gamma-\gamma')/2}.
\end{aligned}$$

Using the Chernoff bound,

$$\Pr[X \geq n^{1-(\gamma-\gamma')/3}] \leq e^{-\Omega(n^{1-(\gamma-\gamma')/6})}.$$

The number of solutions  $V'$  such that  $|V'| < n^{(\gamma+\gamma')/2}$  is at most  $\binom{n}{n^{(\gamma+\gamma')/2}} \leq \left(\frac{en}{n^{(\gamma+\gamma')/2}}\right)^{n^{(\gamma+\gamma')/2}} \leq e^{O(n^{(\gamma+\gamma')/2} \cdot \log n)}$ . Thus, the probability that there is a solution  $V'$  such that  $|V'| < n^{(\gamma+\gamma')/2}$  and  $|W'_p| > O(n^{1-(\gamma-\gamma')/3})$  and  $V', W'_p$  forms a  $(|V'|, |W'_p|)$ -BIS in  $G_p$  is at most

$$(3.4) \quad e^{O(n^{(\gamma+\gamma')/2} \cdot \log n)} \cdot e^{-\Omega(n^{1-(\gamma-\gamma')/6})} \leq e^{-\Omega(n^{1-(\gamma-\gamma')/6})},$$

since  $1 > \frac{2\gamma}{3} + \frac{\gamma'}{3}$ .

Therefore, using (3.3) and (3.4), with high probability over all possible solutions  $V', |W'_p| \leq O(n/\log^{\delta'} n)$  as wanted.  $\square$

**COROLLARY 3.13.** *If  $G$  is a No instance then, with high probability, every solution to unique coverage for  $H$  has size at most  $O(\gamma'' n \log^{1-\delta'} n)$ .*

*Proof.* From Lemma 3.12 and because  $p = \gamma'' \log n$ , it follows that, with high probability, the number of vertices uniquely covered by type-1 edges is at most  $O(\gamma'' n \log^{1-\delta'} n)$ . Combining this bound with Lemma 3.9 shows that, if  $G$  is a No instance (i.e., has no  $(n^{\gamma'}, n/\log^{\delta'} n)$ -BIS), then the size of any solution to unique coverage for  $H$  is at most  $O(\gamma'' n \log^{1-\delta'} n)$ .  $\square$

*Proof.* [of Theorem 3.3] The proof follows easily from Corollaries 3.10 and 3.13 and the assumption that  $BBIS(\gamma, \gamma', \delta, \delta')$  is hard.  $\square$

**3.2. Proving Specific Hardness Results for Unique Coverage.** In this subsection we prove Theorems 3.5, 3.6, and 3.7. In order to prove these theorems, we will prove some hardness results for  $BBIS(\gamma, \gamma', \delta, \delta')$  and then combine them with Theorem 3.3.

Recently, two hardness results for  $BBIS$  were proved by Feige [19] and Khot [32] under different complexity assumptions. Feige [19] proved a constant factor hardness result for  $BBIS$  under the R3SAT hardness hypothesis (for more details see [19]):

**THEOREM 3.14.** [19] *For every  $\varepsilon > 0$  and a given bipartite graph  $G(A \cup B, E)$  with  $|A| = |B| = n$ , deciding between the following two cases is hard, under the complexity assumption that refuting random instances of 3SAT is hard on average:*

1.  $G$  has a  $BBIS$  of size at least  $(\frac{1}{4} - \varepsilon)n$ .
2. Every  $BBIS$  of  $G$  has size smaller than  $(\frac{1}{8} + \varepsilon)n$ .

More recently, Khot [32] proved a similar result, for some (unspecified) constants  $\alpha$  and  $\beta$  instead of  $(\frac{1}{4} - \varepsilon)$  and  $(\frac{1}{8} + \varepsilon)$ , respectively, but under a more plausible assumption that NP problems do not have subexponential-time algorithms. More specifically, he proved the following PCP (probabilistically checkable proof) theorem:

**THEOREM 3.15.** [32] *For every  $\varepsilon > 0$  an integer  $d = O(\frac{1}{\varepsilon} \log(\frac{1}{\varepsilon}))$  exists such that there is a PCP verifier for SAT instances of size  $n$  such that the following hold:*

1. The proof  $\Pi$  for the verifier has size  $2^{n^\varepsilon}$ .
2. The verifier queries a set  $Q$  of size  $d$  bits from  $\Pi$ .
3. Every query bit is uniformly distributed over  $\Pi$ .
4. Completeness: If SAT is a Yes instance,  $\Pi$  is a correct proof, and  $\Pi_0$  is the set of 0-bits in the proof (it contains half the bits from the proof), then

$$\Pr[Q \subseteq \Pi_0] \geq (1 - O(\frac{1}{d})) \frac{1}{2^{d-1}},$$

where the probability is taken over the random tests of the verifier.

5. Soundness: If SAT is a No instance and  $\Pi_*$  is any set of half the bits from  $\Pi$ , then

$$\Pr[Q \subseteq \Pi_*] = \frac{1}{2^d} \pm O(\frac{1}{2^{20d}}).$$

A direct application of Theorem 3.15 implies the following (see [32]):<sup>2</sup>

**THEOREM 3.16.** *Let  $\varepsilon > 0$  be an arbitrary constant and  $d = O(\frac{1}{\varepsilon} \log(\frac{1}{\varepsilon}))$ . Consider an instance  $\Phi$  of SAT with  $n$  variables. Let  $\alpha = \frac{3}{4} \frac{1}{2^{d-1}}$  and  $\beta = (1 - \frac{1}{2^{5d}})\alpha$ . We can construct a bipartite graph  $G(A \cup B, E)$  as an instance of BBIS from the PCP verifier of Theorem 3.15 with  $|A| = |B| = N$  where  $N = 2^{n^\varepsilon}$  such that the following hold:*

- Yes instance: If  $\Phi$  is a Yes instance, then  $G$  has a BBIS of size  $\alpha N$ .
- No instance: If  $\Phi$  is a No instance, then no BBIS of  $G$  has size  $\beta N$ .

**COROLLARY 3.17.** *Assuming that  $\text{NP} \not\subseteq \text{BPTIME}(2^{n^\varepsilon})$ , it is hard to distinguish between the Yes and No cases in the above theorem.*

In order to get a hardness for  $\text{BBIS}(\gamma, \gamma', \delta, \delta')$ , we need a stronger version of Theorem 3.16. For this, we boost the gap in Theorem 3.16 using the standard technique of graph products (see, for example, [20, 11]). Note that Theorem 1.2 in [32] amplifies the gap in Theorem 3.16 using the same technique. However, we require a gap which is asymmetric with respect to the sizes of sets selected on different parts, i.e., the bipartite independent set is not necessarily balanced. In particular, the gap created on one side (say  $A$ ) is polynomial, whereas the gap created on the other side (that is,  $B$ ) is polylogarithmic. Our proof is very similar to that of Theorem 1.2 in [32]. We need the following definition for our proof.

**DEFINITION 3.18.** *For a bipartite graph  $G(A \cup B, E)$  and integers  $K_A, K_B \geq 2$  the bipartite graph  $G^{(K_A, K_B)}$  is defined as follows:*

- The vertex set of  $G^{(K_A, K_B)}$  is  $A' \cup B'$ , where  $A' \cap B' = \emptyset$ ,  $A' = A^{K_A}$ , and  $B' = B^{K_B}$ , i.e.,  $A'$  and  $B'$  are the sets of all  $K_A$ -tuples from  $A$  and all  $K_B$ -tuples from  $B$ , respectively.
- Two vertices  $(a_1, \dots, a_{K_A}) \in A'$  and  $(b_1, \dots, b_{K_B}) \in B'$  are adjacent in  $G^{(K_A, K_B)}$  if and only if  $\exists i, j, 1 \leq i \leq K_A, 1 \leq j \leq K_B, (a_i, b_j) \in E$ .

Suppose that  $G(A \cup B, E)$  is a bipartite graph with  $|A| = |B| = N$ ,  $0 < \alpha < 1$  is a constant, and  $K_A, K_B$  are integers such that

$$(3.5) \quad \frac{1}{\alpha^{K_A}}, \frac{1}{\alpha^{K_B}} \in O(N).$$

Let  $G^*(A^* \cup B^*, E^*)$  be a random subgraph of  $G^{(K_A, K_B)}(A' \cup B', E')$  with  $|A^*| = |B^*| = M$  where  $M = N^3$  and every vertex of  $G^{(K_A, K_B)}$  is selected uniformly at random

<sup>2</sup>Khot defines the bi-clique problem and proves this theorem for bi-clique.

but with different probabilities for  $A'$  and  $B'$ .

LEMMA 3.19. *If  $G(A \cup B, E)$  has a BBIS of size  $\alpha N$  then, with high probability,  $G^*$  has a  $(\frac{1}{2}\alpha^{K_A} M, \frac{1}{2}\alpha^{K_B} M)$ -BIS.*

*Proof.* Let  $A_I \subseteq A$  and  $B_I \subseteq B$  be subsets that form a BBIS of size  $\alpha N$  in  $G$ . Clearly, the subgraph of  $G^{K_A, K_B}$  induced on  $A_I^{K_A} \cup B_I^{K_B}$  is an independent set. Because the vertices of  $G^*$  are selected randomly, each vertex of  $A^*$  belongs to  $A_I^{K_A}$  with probability  $\alpha^{K_A}$ . Also, each vertex of  $B^*$  belongs to  $B_I^{K_B}$  with probability  $\alpha^{K_B}$ . Therefore,  $\mathbb{E}[|A^* \cap A_I^{K_A}|] = \alpha^{K_A} |A^*|$  and  $\mathbb{E}[|B^* \cap B_I^{K_B}|] = \alpha^{K_B} |B^*|$ . Using the Chernoff bound and (3.5),

$$\Pr \left[ |A^* \cap A_I^{K_A}| \leq \frac{1}{2} \alpha^{K_A} |A^*| \right] \leq 2^{-\Omega(N^2)}.$$

Similarly, with high probability,  $|B^* \cap B_I^{K_B}| \geq \frac{1}{2} \alpha^{K_B} |B^*|$ . Therefore, with high probability,  $G^*$  has a  $(\frac{1}{2}\alpha^{K_A} |A^*|, \frac{1}{2}\alpha^{K_B} |B^*|)$ -BIS.  $\square$

LEMMA 3.20. *If  $G(A \cup B, E)$  has no BBIS of size  $\beta N$  then, with high probability,  $G^*$  does not have any  $(2\beta^{K_A} M, 2\beta^{K_B} M)$ -BIS.*

*Proof.* First, note that every maximal bipartite independent set of  $G^{K_A, K_B}$  is of the form  $A_I^{K_A} \cup B_I^{K_B}$ , where  $A_I \cup B_I$  is a bipartite independent set in  $G$ . Consider a fixed maximal bipartite independent set of  $G^{K_A, K_B}$ , say  $A_I^{K_A} \cup B_I^{K_B}$ . Either  $|A_I^{K_A}| < \beta^{K_A} N^{K_A}$  or  $|B_I^{K_B}| < \beta^{K_B} N^{K_B}$ . Without loss of generality, assume  $|A_I^{K_A}| < \beta^{K_A} N^{K_A}$ . Because the element in  $A^*$  and  $B^*$  are selected uniformly randomly:  $\mathbb{E}[|A^* \cap A_I^{K_A}|] < \beta^{K_A} |A^*|$ . Using the Chernoff bound,

$$\Pr \left[ |A^* \cap A_I^{K_A}| \geq 2\beta^{K_A} |A^*| \right] \leq 2^{-\Omega(N^2)}.$$

An almost identical argument applies if  $|B_I^{K_B}| < \beta^{K_B} N^{K_B}$ . Because there are at most  $2^{O(N)}$  possible maximal bipartite independent sets in  $G$ , using union bound, the probability of having a  $(2\beta^{K_A} |A^*|, 2\beta^{K_B} |B^*|)$ -BIS in  $G^*$  is  $o(1)$ .  $\square$

Let  $\Phi$  be an instance of SAT and let  $\varepsilon > 0$  be an arbitrary small constant. Define  $d$ ,  $\alpha$ ,  $\beta$ , and  $G(A \cup B, E)$  as in Theorem 3.16, with  $|A| = |B| = N$ . Also let  $M = N^3$ ,  $K_A = -\frac{(1-\gamma)\log M}{\log \alpha}$ , and  $K_B = -\frac{\delta \log \log M}{\log \alpha}$ , for some constants  $0 < \gamma, \delta < 1$  such that  $1/\alpha^{K_A} \in O(N)$ . Construct the graph  $G^{K_A, K_B}$  and the random subgraph of it  $G^*(A^* \cup B^*, E^*)$ , where  $|A^*| = |B^*| = M$ , as explained above. By Theorem 3.16 and Lemmas 3.19 and 3.20, we have the following:

1. If  $\Phi$  is a Yes instance then, by Theorem 3.16,  $G$  has a BBIS of size  $\alpha N$ . So, by Lemma 3.19, with high probability,  $G^*$  has a  $(\frac{1}{2}\alpha^{K_A} M, \frac{1}{2}\alpha^{K_B} M)$ -BIS. By definition of  $K_A$  and  $K_B$ , this is a  $(\frac{M^\gamma}{2}, \frac{M}{2 \log^\delta M})$ -BIS in  $G^*$ .
2. If  $\Phi$  is a No instance then, by Theorem 3.16,  $G$  has no BBIS of size  $\beta N$ . So, by Lemma 3.20, with high probability,  $G^*$  has no  $(2\beta^{K_A} M, 2\beta^{K_B} M)$ -BIS. With  $\ell = \log_\alpha(\beta/\alpha)$ ,  $\gamma' = \gamma - \ell(1 - \gamma)$ , and  $\delta' = \delta(1 + \ell)$ , this means that, with high probability,  $G^*$  has no  $(2M^{\gamma'}, \frac{2M}{\log^{\delta'} M})$ -BIS.

Therefore, we have proved the following amplified version of Theorem 3.16.

THEOREM 3.21. *Let  $G(A \cup B, E)$  be a given bipartite graph with  $|A| = |B| = n$ , together with an arbitrary small constant  $\varepsilon > 0$ , and  $d = O(\frac{1}{\varepsilon} \log(\frac{1}{\varepsilon}))$ ,  $\alpha = \frac{3}{4} \frac{1}{2^d - 1}$ ,  $\beta = (1 - \frac{1}{2^{5d}})\alpha$ . Furthermore let  $0 < \gamma' < \gamma \leq 1$  and  $0 \leq \delta < \delta' \leq 1$  be such that  $\delta$  is any constant and with  $\ell = \log_\alpha(\beta/\alpha)$ :  $\gamma' = \gamma - \ell(1 - \gamma)$ , and  $\delta' = \delta(1 + \ell)$ . Then it is hard to distinguish between the following two cases unless  $\text{NP} \subseteq \text{BPTIME}(2^{n^\varepsilon})$ :*

1.  $G$  has a  $(\frac{n^\gamma}{2}, \frac{n}{2 \log^\delta n})$ -BIS.
2.  $G$  has no  $(2n^{\gamma'}, \frac{2n}{\log^{\delta'} n})$ -BIS.

*Proof.* [of Theorem 3.5] Given a bipartite graph  $G(A \cup B, E)$  with  $|A| = |B| = n$  as an instance of bipartite independent set and parameters  $\varepsilon, \alpha, \beta, \delta, \gamma, \delta',$  and  $\gamma'$  as in Theorem 3.21 we construct  $H(V \cup W, F)$  as explained in the proof of Theorem 3.3. We choose  $\delta = \frac{1}{1+\ell}$ , where  $\ell = \log_\alpha(\beta/\alpha)$ . Therefore  $\delta' = 1$  and by Corollaries 3.10 and 3.13 unless  $\text{NP} \subseteq \text{BPTIME}(2^{n^\varepsilon})$  it is hard to approximate unique coverage within a factor of  $\Omega(1/\log^{\delta'-\delta} n)$ . Because  $\delta' - \delta = \frac{\ell}{1+\ell}$  and  $\ell$  is a function of  $\varepsilon$ , this completes the proof of the theorem.  $\square$

*Proof.* [of Theorem 3.6] If our starting point to prove Theorem 3.21 is Theorem 3.14 instead of Theorem 3.16, then we have  $\alpha = \frac{1}{4} - \varepsilon$  and  $\beta = \frac{1}{8} + \varepsilon$ , and  $\ell = \frac{1}{2} + \varepsilon'$ , where  $\varepsilon' = \varepsilon'(\varepsilon)$  is a constant. Then the same argument as in the proof of Theorem 3.5 proves a hardness of  $O(1/\log^{\frac{1/2+\varepsilon'}{1+1/2+\varepsilon'}} n)$  which is  $O(1/\log^{\frac{1}{3}-\sigma} n)$  for some  $\sigma = \sigma(\varepsilon)$ .  $\square$

We now turn to the proof of Theorem 3.7. It is based on the following hypothesis.

**HYPOTHESIS 3.22.** *Given a bipartite graph  $G(A \cup B, E)$  with size  $|A| = |B| = n$  as an instance of BBIS and for absolute constants  $1 \geq \gamma > \gamma' > 0$  it is hard to distinguish the following two cases:*

1.  $G$  has an  $(n^\gamma, \Omega(n))$ -BIS.
2.  $G$  has no  $(n^{\gamma'}, n/\log n)$ -BIS.

Now we show how Hypothesis 3.22 would imply an  $O(1/\log n)$ -hardness for unique coverage.

*Proof.* [of Theorem 3.7] Given a bipartite graph  $G(A \cup B, E)$  with size  $|A| = |B| = n$  and  $1 \geq \gamma > \gamma' > 0$  we construct  $H(V \cup W, F)$ , the instance of unique coverage, as in the reduction of Theorem 3.3.

- If  $G$  has an  $(n^\gamma, \Omega(n))$ -BIS, then  $H$  has a unique coverage of size  $\Omega((\gamma - \gamma')n \log^{1-\delta} n)$  with  $\delta = 0$ , which is  $\Omega(n \log n)$ .
- If  $G$  has no  $(n^{\gamma'}, n/\log n)$ -BIS, then every unique coverage solution for  $H$  has size at most  $O((\gamma - \gamma')n \log^{1-\delta'} n)$  with  $\delta' = 1$ , which is  $O(n)$ .

This implies that, assuming Hypothesis 3.22, it is hard to distinguish between the two cases above, and hence hard to approximate unique coverage within a factor of  $\Omega(1/\log n)$ .  $\square$

The authors suspect that Hypothesis 3.22 will be difficult to refute in the near future. The BBIS problem appears to be at least as hard to approximate as maximum independent set in general graphs. (This is not a theorem, but merely an empirical observation concerning currently known approximation algorithms.) For the latter problem, despite extensive work, no known polynomial-time algorithm can distinguish between graphs with independent sets of size  $\Omega(n/k)$  and graphs with no independent set of size  $n^{1/k}$ , where  $k$  is some sufficiently large constant. It is plausible (though not certain) that any refutation of Hypothesis 3.22 would lead to major improvements in the approximation ratio for maximum independent sets in general graphs.

## 4. Approximation Algorithms.

**4.1.  $\Omega(1/\log m)$ -Approximation.** In this section we develop our main logarithmic approximation algorithm:

**THEOREM 4.1.** *There is an  $\Omega(1/\log \rho) = \Omega(1/\log m)$ -approximation algorithm for the budgeted unique coverage problem, where  $\rho$  is one more than the ratio of the*

maximum number of sets in which an element appears over the minimum number of sets in which an element appears.

*Proof.* First we find an  $(1-1/e)$ -approximate solution  $\mathcal{S}'$  to the maximum coverage problem with the same universe, profits, sets, costs, and budget [33]. Because the total profit of uniquely covered elements is always at most the total profit of all covered elements, the optimum solution value OPT to the unique coverage problem must be at most the optimum solution value to the maximum coverage problem. Thus the total profit of covered elements in  $\mathcal{S}'$  is within an  $1 - 1/e$  factor of an upper bound on OPT. Symbolically, if  $p(S)$  denotes the total profit of elements in set  $S$  and  $\bigcup \mathcal{S}'$  denotes the union  $\bigcup_{S \in \mathcal{S}'} S$ , then  $p(\bigcup \mathcal{S}') \geq (1 - 1/e) \text{OPT}$ .

We cluster the elements in  $\bigcup \mathcal{S}'$  into  $\lg \rho$  groups as follows: an element is in *group*  $i$  if it is covered by between  $2^i$  and  $2^{i+1} - 1$  sets. The group  $i^*$  with the most total profit must have at least a  $1/\lg \rho$  fraction of  $p(\bigcup \mathcal{S}') \geq (1 - 1/e) \text{OPT}$ . Now we randomly discard sets from  $\mathcal{S}'$ , keeping a set with probability  $1/2^{i^*}$ . We claim that, in expectation, the resulting collection  $\mathcal{S}''$  uniquely covers a constant fraction of the elements in group  $i^*$ , which is  $\Omega(\text{OPT} / \log \rho)$ .

Fix an element  $x$  in group  $i^*$ , and suppose that it was covered  $d$  times in  $\mathcal{S}'$ ,  $2^{i^*} \leq d \leq 2^{i^*+1} - 1$ . The probability that  $x$  is covered exactly once by  $\mathcal{S}''$  is  $(d/2^{i^*})(1 - 1/2^{i^*})^{d-1}$ . (There is a factor of  $d$  for the choice of which set covers  $x$ , a  $1/2^{i^*}$  probability that this set is kept, and a  $1 - 1/2^{i^*}$  probability that each of the  $d - 1$  other sets is discarded.) By our bounds on  $d$ , the probability that  $x$  is covered exactly once by  $\mathcal{S}''$  is at least  $(1 - 1/2^{i^*})2^{i^*+1} \geq 1/e^2$ .

The expected total profit of elements covered exactly once by  $\mathcal{S}''$  is at least  $\sum \{p_x/e^2 \mid x \text{ in group } i^*\}$ , which is  $1/e^2$  times the total profit of elements in group  $i^*$ , which we argued is at least  $(1 - 1/e) \text{OPT} / \lg \rho$ . Therefore the expected profit of our randomized solution is at least  $(1/e^2 - 1/e^3) \text{OPT} / \lg \rho = \Omega(\text{OPT} / \log \rho)$ .

We can derandomize this algorithm by the standard method of conditional expectation [44]. For each set in  $\mathcal{S}'$ , we decide whether to keep it in  $\mathcal{S}''$  by trying both options, and choosing the option that maximizes the conditional expectation of the total profit of elements in group  $i^*$  uniquely covered by  $\mathcal{S}''$ . The conditional expectations can be computed easily in polynomial time according to the analysis above.  $\square$

The approximate solution computed by this algorithm is not only within an  $\Omega(1/\log m)$  factor of the optimal unique coverage, but also within an  $\Omega(1/\log m)$  of the optimal maximum coverage. As a consequence, we also obtain an  $\Omega(1/\log m)$ -approximation for the more general problem of budgeted low-coverage described in Section 2.1.

**4.2. Approximation with Bounded Set Size.** In this section we consider the unique coverage problem with a bound  $B$  on the maximum set size, or more generally, the budgeted unique coverage problem with a bound  $B$  on the ratio between the maximum profit of a set and the minimum profit of an element (recall that the profit of a set is the sum of profits of its elements). In both cases we obtain an approximation ratio of  $\Omega(1/\log B)$ . In particular,  $B \leq n$ , so this algorithm is an  $\Omega(1/\log n)$ -approximation.

**THEOREM 4.2.** *There is an  $\Omega(1/\log B)$ -approximation algorithm for instances of the budgeted unique coverage problem in which the minimum element profit is 1 and the total profit of every set is at most  $B$ .*

*Proof.* As before, we first find an  $(1 - 1/e)$ -approximate solution  $\mathcal{S}'$  to the maximum coverage problem with the same universe, profits, sets, costs, and budget [33].



As argued in the proof of Theorem 4.1,  $p(\bigcup S') \geq (1 - 1/e) \text{OPT}$ , where  $p(S)$  denotes the total profit of elements in set  $S$ ,  $\bigcup S'$  denotes the union  $\bigcup_{S \in \mathcal{S}'} S$ , and  $\text{OPT}$  denotes the optimum solution value to the unique coverage problem.

We modify  $\mathcal{S}'$  to be minimal by removing any sets that do not uniquely cover any elements. Thus the set of covered elements remains the same, so the same upper bound on  $\text{OPT}$  holds. Let  $X$  be the set of elements covered by exactly one set of  $\mathcal{S}'$ . Because  $\mathcal{S}'$  is minimal, each set must uniquely cover at least one element in  $X$ , so  $|X| \geq |\mathcal{S}'|$ . Because every element has profit at least 1,  $p(X) \geq |X| \geq |\mathcal{S}'|$ .

If  $p(\bigcup S') \leq 2|\mathcal{S}'| \leq 2p(X)$ , then  $\mathcal{S}'$  is already an  $\Omega(1)$ -approximate solution to the budgeted unique coverage problem. If  $p(\bigcup S') > 2|\mathcal{S}'|$ , then we claim that the total profit of elements covered at most  $B$  times by  $\mathcal{S}'$  is at least  $p(\bigcup S')/2$ . Otherwise, the elements covered more than  $B$  times by  $\mathcal{S}'$  would be at least  $p(\bigcup S')/2$ , and thus the total profit of the sets would satisfy  $\sum_{S \in \mathcal{S}'} p(S) > Bp(\bigcup S')/2 > B|\mathcal{S}'|$ , contradicting that every set in  $\mathcal{S}$  (and thus  $\mathcal{S}'$ ) has total profit at most  $B$ . Now we apply Theorem 4.1 above to the elements covered at most  $B$  times by  $\mathcal{S}'$ , for which  $\rho \leq B$ . Thus we obtain an  $\Omega(1/\log B)$ -approximation for this subproblem, whose optimal solution value is at least  $(1 - 1/e) \text{OPT} / 2$ .  $\square$

We note that, when every set has cardinality at most  $B = 3$  and every element appears in exactly two sets ( $\rho = 1$ ), the unique coverage problem is exactly max-cut in maximum-degree-3 graphs, so the problem is APX-hard even in this restricted case [46, 2].

**Appendix A. Randomized Rounding for Envy-Free Pricing.** In this section we prove the necessary lemma about randomized rounding needed in Section 2.2 for the reduction from unique coverage to unlimited-supply single-minded envy-free pricing.

**LEMMA A.1.** *In the setting of single-minded envy-free pricing, suppose all valuations are 1. Then there is a price assignment that uses prices of just 0 and 1 and whose profit is within a constant factor of optimal.*

*Proof.* Consider the optimal assignment of prices  $p_i$  to items  $I_i$ . If any price  $p_i$  is larger than 1, we set it to 1 at no cost. Now we round by setting the new price  $p'_i$  of item  $I_i$  to 1 with probability  $\frac{1}{2}p_i$  and to 0 otherwise. We claim that, if  $u_i = \sum_{I_j \in B_i} p_j < 1$  (i.e., the optimal solution profits  $u_i$  from buyer  $b_i$ ), then the probability that the seller profits 1 from buyer  $b_i$  is at least  $\frac{1}{2e}u_i$ .

The probability that the seller profits 1 from buyer  $b_i$ , who desires bundle  $B_i$ , is  $\sum_{I_j \in B_i} \frac{1}{2}p_j \prod_{I_k \neq I_j \in B_i} (1 - \frac{1}{2}p_k)$ . This quantity can be rewritten as  $\prod_{I_k \in B_i} (1 - \frac{1}{2}p_k) \sum_{I_j \in B_i} \frac{1}{2}p_j / (1 - \frac{1}{2}p_j)$ . Because  $\sum_{I_j \in B_i} p_j \leq \frac{1}{2}$ , it is easy to show that the quantity is minimized when all of the  $p_j$ 's,  $I_j \in B_i$ , are equal. Thus the probability of profit from  $b_i$  is at least  $(1 - \frac{1}{2}u_i/|B_i|)^{|B_i|} \frac{1}{2}u_i / (1 - \frac{1}{2}u_i/|B_i|)$ . Because  $1 - x \geq e^{-2x}$  for  $0 \leq x \leq \frac{1}{2}$ , this probability is at least  $e^{-u_i} \frac{1}{2}u_i \geq e^{-1} \frac{1}{2}u_i$ , as claimed.

Thus the expected total profit in the modified solution is at least  $\sum_i \frac{1}{2e}u_i/e$ , which is  $\frac{1}{2e}$  times the profit of the optimal solution. We can derandomize this algorithm by the standard method of conditional expectation [44]; see the proof of Theorem 4.1.  $\square$

**Acknowledgments.** We thank David Abusch-Magder, who mentioned to us a real-world application in wireless networks at Bell Labs, of which unique coverage is a natural simplification. We also thank David Karger for initial discussions about the problem. The third author also thanks Nikhil Bansal, Venkatesan Guruswami, Jason Hartline, Kamal Jain, and Kunal Talwar for fruitful discussions. Finally, we thank the anonymous referees for their helpful comments.

## REFERENCES

- [1] G. AGGARWAL, T. FEDER, R. MOTWANI, AND A. ZHU, *Algorithms for multi-product pricing*, in Proceedings of the 31st International Colloquium on Automata, Languages and Programming, vol. 3142 of Lecture Notes in Computer Science, 2004, pp. 72–83.
- [2] P. ALIMONTI AND V. KANN, *Hardness of approximating problems on cubic graphs*, in Proceedings of the 3rd Italian Conference on Algorithms and Complexity, vol. 1203 of Lecture Notes in Computer Science, 1997, pp. 288–298.
- [3] N. ALON, A. BAR-NOY, N. LINIAL, AND D. PELEG, *A lower bound for radio broadcast*, Journal of Computer and System Sciences, 43 (1991), pp. 290–298.
- [4] C. AMBÜHL, A. E. F. CLEMENTI, M. D. IANNI, N. LEV-TOV, A. MONTI, D. PELEG, G. ROSSI, AND R. SILVESTRI, *Efficient algorithms for low-energy bounded-hop broadcast in ad-hoc wireless networks*, in Proceedings of the 21st Annual Symposium on Theoretical Aspects of Computer Science, vol. 2996 of Lecture Notes in Computer Science, February 2004, pp. 418–427.
- [5] M. ANDREWS, J. CHUZHUY, S. KHANNA, AND L. ZHANG, *Hardness of the undirected edge-disjoint paths problem with congestion*, in Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, 2005, pp. 226–244.
- [6] M. ANDREWS AND L. ZHANG, *Hardness of the undirected edge-disjoint paths problem*, in Proceedings of the 37th Annual ACM Symposium on Theory of Computing, 2005, pp. 276–283.
- [7] A. ARCHER, C. PAPADIMITRIOU, K. TALWAR, AND É. TARDOS, *An approximate truthful mechanism for combinatorial auctions with single parameter agents*, in Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms, 2003, pp. 205–214.
- [8] V. BAHL, M. HAJIAGHAYI, K. JAIN, V. S. MIRROKNI, L. QUI, AND A. SABERI, *Cell breathing in wireless lan: Algorithms and evaluation*. Manuscript, 2005.
- [9] N. BANSAL, A. BLUM, S. CHAWLA, AND A. MEYERSON, *Approximation algorithms for deadline-TSP and vehicle routing with time-windows*, in Proceedings of the 36th Annual ACM Symposium on Theory of Computing, 2004, pp. 166–174.
- [10] R. BAR-YEHUDA, O. GOLDREICH, AND A. ITAI, *On the time-complexity of broadcast in multi-hop radio networks: an exponential gap between determinism and randomization*, Journal of Computer and System Sciences, 45 (1992), pp. 104–126.
- [11] P. BERMAN AND G. SCHNITGER, *On the complexity of approximating the independent set problem*, Information and Computation, 96 (1992), pp. 77–94.
- [12] I. CHLAMTAC AND O. WEINSTEIN, *The wave expansion approach to broadcasting in multi-hop radio networks*, IEEE Transactions on Communications, 39 (1991), pp. 426–433.
- [13] J. CHUZHUY, S. GUHA, E. HALPERI, S. KHANNA, G. KORTSARZ, AND J. S. NAOR, *Asymmetric  $k$ -center is  $\log^* n$ -hard to approximate*, in Proceedings of the 36th Annual ACM Symposium on Theory of Computing, 2004, pp. 21–27.
- [14] X. DENG, C. PAPADIMITRIOU, AND S. SAFRA, *On the complexity of equilibria*, in Proceedings of the 34th Annual ACM Symposium on Theory of Computing, 2002, pp. 67–71.
- [15] N. R. DEVANUR, C. H. PAPADIMITRIOU, A. SABERI, AND V. V. VAZIRANI, *Market equilibrium via a primal-dual-type algorithm*, in Proceedings of the 43rd Symposium on Foundations of Computer Science, 2002, pp. 389–395.
- [16] M. ELKIN AND G. KORTSARZ, *Combinatorial logarithmic approximation algorithm for directed telephone broadcast problem*, in Proceedings of the 34th Annual ACM Symposium on Theory of Computing, 2002, pp. 438–447.
- [17] ———, *Polylogarithmic inapproximability of the radio broadcast problem*, in Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, 2004, pp. 105–116.
- [18] U. FEIGE, *A threshold of  $\ln n$  for approximating set cover*, Journal of the ACM, 45 (1998), pp. 634–652.
- [19] ———, *Relations between average case complexity and approximation complexity*, in Proceedings of the 34th Annual ACM Symposium on Theory of Computing, 2002, pp. 534–543.
- [20] U. FEIGE, S. GOLDWASSER, L. LOVÁSZ, S. SAFRA, AND M. SZEGEDY, *Interactive proofs and the hardness of approximating cliques*, Journal of the ACM, 43 (1996), pp. 268–292.
- [21] U. FEIGE, M. M. HALLDÓRSSON, G. KORTSARZ, AND A. SRINIVASAN, *Approximating the domatic number*, SIAM Journal on Computing, 32 (2002), pp. 172–195.
- [22] Y. FUKUDA, T. ABE, AND Y. OIE, *Decentralized access point selection architecture for wireless LANs: Deployability and robustness for wireless LANs*, in In Proc. of IEEE VTC, September 2004.
- [23] L. GASNIENEC, D. PELEG, AND Q. XIN, *Faster communication in known topology radio networks*, in Proceedings of the 24th Annual ACM SIGACT-SIGOPS Symposium on Principles of

- Distributed Computing, 2005, pp. 129–137.
- [24] M. X. GOEMANS AND D. P. WILLIAMSON, *Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming*, Journal of the ACM, 42 (1995), pp. 1115–1145.
- [25] S. GUHA AND S. KHULLER, *Greedy strikes back: improved facility location algorithms*, Journal of Algorithms, 31 (1999), pp. 228–248.
- [26] F. GUL AND E. STACCHETTI, *Walrasian equilibrium with gross substitutes*, Journal of Economic Theory, 87 (1999), pp. 95–124.
- [27] V. GURUSWAMI, J. D. HARTLINE, A. R. KARLIN, D. KEMPE, C. KENYON, , AND F. MCSHERRY, *On profit-maximizing envy-free pricing*, in Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, 2005, pp. 1164–1173.
- [28] V. GURUSWAMI AND L. TREVISAN, *The complexity of making unique choices: Approximating 1-in-k SAT*, in Proceedings of 8th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, vol. 3624 of Lecture Notes in Computer Science, 2005, pp. 99–110.
- [29] J. HÅSTAD, *Some optimal inapproximability results*, in Proceedings of the 29th Annual ACM Symposium on Theory of Computing, 1997, pp. 1–10.
- [30] D. S. HOCHBAUM, *Approximating covering and packing problems: set cover, vertex cover, independent set, and related problems*, in Approximation algorithms for NP-hard problems, PWS Publishing Co., Boston, 1997, pp. 94–143.
- [31] G. JUDD AND P. STEENKISTE, *Fixing 802.11 access point selection*, in ACM SIGCOMM Poster, August 2002.
- [32] S. KHOT, *Ruling out PTAS for graph min-bisection, densest subgraph and bipartite clique*, in Proceedings of the 45th Annual Symposium on Foundations of Computer Science, 2004, pp. 136–145.
- [33] S. KHULLER, A. MOSS, AND J. S. NAOR, *The budgeted maximum coverage problem*, Information Processing Letters, 70 (1999), pp. 39–45.
- [34] S. KNOT, G. KINDLER, E. MOSSEL, AND R. O’DONNELL, *Optimal inapproximability results for MAX-CUT and other 2-variable CSPs?*, SIAM Journal of Computing, 37 (2007), pp. 319–367.
- [35] D. R. KOWALSKI AND A. PELC, *Broadcasting in undirected ad hoc radio networks*, in Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing, 2003, pp. 73–82.
- [36] ———, *Centralized deterministic broadcasting in undirected multi-hop radio networks*, in Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, 2004, pp. 171–182.
- [37] ———, *Time of deterministic broadcasting in radio networks with local knowledge*, SIAM Journal on Computing, 33 (2004), pp. 870–891.
- [38] M. KRIVELEVICH, Z. NUTOV, M. R. SALAVATIPOUR, J. V. YUSTER, AND R. YUSTER, *Approximation algorithms and hardness results for cycle packing problems*, ACM Transactions on Algorithms, 3 (2007), pp. 48:1–21.
- [39] E. KUSHILEVITZ AND Y. MANSOUR, *An  $\Omega(D \log(N/D))$  lower bound for broadcast in radio networks*, SIAM Journal on Computing, 27 (1998), pp. 702–712.
- [40] D. LEHMANN, L. I. O’CALLAGHAN, AND Y. SHOHAM, *Truth revelation in approximately efficient combinatorial auctions*, in Proceedings of the 1st ACM Conference on Electronic Commerce, 1999, pp. 96–102.
- [41] N. LEV-TOV AND D. PELEG, *Exact algorithms and approximation schemes for base station placement problems*, in Proceedings of the 8th Scandinavian Workshop on Algorithm Theory, vol. 2368 of Lecture Notes in Computer Science, 2002, pp. 90–99.
- [42] C. LUND AND M. YANNAKAKIS, *On the hardness of approximating minimization problems*, Journal of the ACM, 41 (1994), pp. 960–981.
- [43] A. MOSS AND Y. RABANI, *Approximation algorithms for constrained for constrained node weighted steiner tree problems*, in Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, 2001, pp. 373–382.
- [44] R. MOTWANI AND P. RAGHAVAN, *Randomized Algorithms*, Cambridge University Press, 1995.
- [45] A. MU’ALEM AND N. NISAN, *Truthful approximation mechanisms for restricted combinatorial auctions: extended abstract*, in 18th National Conference on Artificial Intelligence, 2002, pp. 379–384.
- [46] C. H. PAPANIMITRIOU AND M. YANNAKAKIS, *Optimization, approximation, and complexity classes*, Journal of Computer and System Sciences, 43 (1991), pp. 425–440.
- [47] R. RAZ AND S. SAFRA, *A sub-constant error-probability low-degree test, and a sub-constant error-probability pcg characterization of np*, in Proceedings of the 29th Annual ACM Sym-

- posium on Theory of Computing, 1997, pp. 475–484.
- [48] M. R. SALAVATIPOUR AND J. VERSTRAETE, *Disjoint cycles: Integrality gap, hardness, and approximation*, in Proceedings of the 11th Conference on Integer Programming and Combinatorial Optimization, 2005, pp. 51–65.
- [49] L. WALRAS, *Elements of Pure Economics*, Allen and Unwin, 1954.