

On Capacitated Set Cover Problems

Nikhil Bansal¹, Ravishankar Krishnaswamy^{2*}, and Barna Saha^{3*}

¹ IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA.

² Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

³ Computer Science Department, University of Maryland College Park, College Park, MD 20740, USA.

Abstract. Recently, Chakrabarty et al. [5] initiated a systematic study of capacitated set cover problems, and considered the question of how their approximability relates to that of the uncapacitated problem on the same underlying set system. Here, we investigate this connection further and give several results, both positive and negative. In particular, we show that if the underlying set system satisfies a certain *hereditary property*, then the approximability of the capacitated problem is closely related to that of the uncapacitated version. We also give related lower bounds, and show that the hereditary property is necessary to obtain non-trivial results. Finally, we give some results for capacitated covering problems on set systems with low hereditary discrepancy and low VC dimension.

1 Introduction

In this paper, we consider the approximability of *capacitated* set cover problems (CSC). In a typical (uncapacitated) set cover instance, we are given a universe X of n elements and a collection \mathcal{S} of m subsets of X , each subset with an associated cost; the goal is to pick the collection of sets $\mathcal{S}' \subseteq \mathcal{S}$ of least total cost, such that each element $e \in X$ is contained in at least one set $S \in \mathcal{S}'$. It is well known that the greedy algorithm for set cover achieves an approximation ratio of $\ln n$, and that in general this approximation factor cannot be improved up to lower order terms [9]. However, in several cases of interest, improved approximation guarantees or even exact algorithms can be obtained. Typical examples are problems arising in network design where the underlying set system may be totally unimodular or have other interesting structural properties [12], or in geometric settings where the set system may have low structural complexity, often measured in terms of VC dimension [3] or union complexity [13]. In general, the study of covering problems is an extensive area of research in both combinatorial optimization and algorithms.

In the capacitated version of the set cover problem, the elements additionally have *demands* $d : X \rightarrow \mathbb{R}^+$, sets have *supplies* $s : \mathcal{S} \rightarrow \mathbb{R}^+$, and the goal is to find a minimum cost collection of sets \mathcal{S}' such that for each element e , the total supply of sets in \mathcal{S}' that cover e is at least $d(e)$. A general CSC is defined by the following integer linear program:

* Work done while visiting IBM T. J. Watson Research Yorktown Heights.

$$\begin{aligned} \text{CSC}(A, d, s, c) \quad \min \quad & \sum_{i=1}^m c_i x_i & (1) \\ \text{s.t.} \quad & \sum_{i=1}^m A_{ij} s_i x_i \geq d_j & \forall 1 \leq j \leq n & (2) \\ & x_i \in \{0, 1\} & \forall S_i \in \mathcal{S} & (3) \end{aligned}$$

Here, s_i denotes the supply of set S_i , d_j is the demand of element $j \in [n]$, and A is the $\{0, 1\}$ incidence matrix of the set system. The variable x_i indicates whether S_i is chosen or not, and hence the constraints ensure that for each element $j \in [n]$, the total supply of sets containing it is at least d_j .

Capacitated covering problems arise naturally in a variety of scenarios. For example, consider the minimum Steiner tree problem where the goal is to find the minimum-cost subgraph connecting terminals to a root. This can be cast as a set cover problem, viewing each graph cut separating some terminal from the root as an element, and each edge in the graph as a set (that covers every cut that it crosses). Now, if the terminals have a bandwidth requirement, and the edges have different bandwidth capacities, this corresponds to a capacitated covering problem. Similar generalizations naturally arise for most uncapacitated covering problems. Capacitated covering problems also arise indirectly as subroutines in other problems. For example, Bansal and Pruhs [1] showed that the scheduling problem of minimizing arbitrary functions of flow time on a single machine is equivalent (up to $O(1)$ factors) to the capacitated version of the geometric set cover problem of covering points in \mathbb{R}^2 using axis-aligned rectangles all of which touch the x -axis.

While capacitated covering problems have been studied previously, Chakrabarty et al. [5] recently initiated a more systematic study of these problems. Motivated by the extensive existing works on the uncapacitated set cover problem, they considered the following natural question. *Is there a relationship between the approximability of a capacitated set cover problem and the uncapacitated problem on the same underlying set system?* In particular, is it possible to exploit the combinatorial structure of the underlying incidence matrix in the set cover problem to design good algorithms for the capacitated case?

To understand this question better, it is instructive to even consider the case of the simplest possible set system: that with a *single* element. In this case, the problem reduces to precisely the so-called Knapsack Cover problem, where given a knapsack (element) of demand B and items (sets) with supplies s_1, \dots, s_m and costs c_1, \dots, c_m , the goal is to find a minimum cost collection of items that covers the knapsack. Already here, it turns out that the natural LP relaxation⁴ of the integer program (1)-(3) has arbitrarily large integrality gap.⁵ In a celebrated result, Carr et al. [4] showed that this natural LP can be strengthened by adding exponentially many so-called Knapsack Cover (KC) inequalities. These inequalities can be separated in polynomial time and hence the LP can be solved efficiently to within any accuracy using the Ellipsoid method. The

⁴ Where we replace the $x \in \{0, 1\}$ in the IP by $x \in [0, 1]$.

⁵ Consider an instance with two items of size $B - 1$ each and costs 0 and 1 respectively. Clearly any integral solution must choose both items, incurring a cost of 1. The LP can however choose the 0 cost item completely, and cost 1 item to extent $1/(B - 1)$, incurring a cost of $1/(B - 1)$.

integrality gap of this strengthened LP reduces to 2, and this is also tight. We remark that there is also a local ratio based interpretation of these KC inequalities [2].

Interestingly, Chakrabarty et al. [5] showed (see Theorem 1 for a formal statement) that given any CSC problem, the natural LP relaxation strengthened by adding KC inequalities for each element has an integrality gap that is no worse (up to $O(1)$ factors) than the integrality gap for two related *uncapacitated* problems. The first of these problems is simply the multi-cover problem on the same set system A , and the second one is the so-called *priority* set cover version of A , that we next define. Thus, roughly speaking their result shows that KC inequalities allow us to forget about capacities, at the expense of somewhat complicating the underlying set system.

Priority Covering Problems: Given a set cover instance specified by the incidence matrix A (the representation could be implicit as in network design problems), a *priority* version of the covering problem (PSC) is defined as follows. The elements and sets have priorities $\pi : X \cup \mathcal{S} \rightarrow \mathbb{Z}^+$. The goal is to pick a minimum cost collection of sets \mathcal{S}' such that for each element j , there is at least one set $S_i \in \mathcal{S}'$ containing j and with priority at least that of j , i.e., $\pi(S_i) \geq \pi(e_j)$.

The natural integer programming formulation for PSC is:

$$\begin{aligned} \text{PSC}(A, \pi, c) \text{ minimize } & \sum_{i=1}^m c_i x_i \\ \text{subject to } & (1) \sum_{i=1}^m A_{ij} \mathbf{1}_{(\pi(S_i) \geq \pi(e_j))} x_i \geq 1 \quad \forall 1 \leq j \leq n \\ & (2) \quad x_i \in \{0, 1\} \quad \forall S_i \in \mathcal{S} \end{aligned}$$

Here, $\mathbf{1}_{(a \geq b)}$ is the indicator variable for the condition inside (i.e., 1 if $a \geq b$ and 0 otherwise). Thus a priority cover problem is an (uncapacitated) set cover problem, with the incidence matrix $B_{ij} = A_{ij} \cdot \mathbf{1}_{(\pi_i \geq \pi_j)}$ instead of A . The structure of B has an interesting geometric connection to that of A . In particular, permute the columns of A in non-decreasing order of supply priorities and rows of A in non-decreasing order of demand priorities. Then, the priority matrix Π defined as $\Pi_{ij} = \mathbf{1}_{(\pi(S_i) \geq \pi(e_j))}$ has a ‘‘stair-case’’ structure of 1’s (see Figure 1 for an illustration), and $B = A \circ \Pi$ is the element-wise product of A and Π . The number of stairs in Π is equal to number of distinct priorities k (which plays an important role in our results later).

Formally, Chakrabarty et al. [5] showed the following result.

Theorem 1 ([5]). *Let $\text{CSC}(A, d, s, c)$ be a capacitated set cover problem instance. Let $\text{MSC}(A, d', \mathbf{1}, c)$ denote the uncapacitated multi-cover⁶ problem with incidence matrix A and covering requirements d' , and let $\text{PSC}(A, \pi, c)$ denote the priority covering problem with incidence matrix A and priorities π . If*

1. *The integrality gap of the natural LP relaxation of $\text{MSC}(A, d', \mathbf{1}, c)$ is at most α for all possible covering requirements d' , and*
2. *The integrality gap of the priority problem $\text{PSC}(A, \pi, c)$ is at most β for all priority functions π ,*

⁶ By multi-cover we mean the usual generalization of standard set cover where an element j may wish to be covered by d_j distinct sets, instead of just one.




Fig. 1. A is an arbitrary $\{0, 1\}$ matrix, $\Pi_{ij} = 1$ if $\pi(S_i) \geq \pi(e_j)$, and $B = A \wedge \Pi$.

Then the integrality gap of the LP relaxation of the capacitated problem $\text{CSC}(A, d, s, c)$ strengthened by KC inequalities is $O(\alpha + \beta)$. Moreover, the number of distinct priorities in the instance $\text{PSC}(A, \pi, c)$ is at most $\log s_{\max}$ where $s_{\max} = \max_{i \in [m]} s_i$ denotes the maximum supply of a set.

Here, as usual, we say that a problem has integrality gap α if for every feasible fractional solution x , there is a feasible integer solution \tilde{x} with cost at most α times the fractional cost. Also note that in the priority cover problem, only the relative values of priorities matter, hence we can assume that the priorities are always integers $1, \dots, k$.

In light of Theorem 1, it suffices to bound the integrality gap of the multi-cover version and priority cover version of the underlying set cover problem. Typically, the multi-cover version is not much harder than the set cover problem itself (e.g. if the matrix is totally unimodular, for various geometric systems, and so on), and the hard work lies in analyzing the priority problem.

We note here that a converse of Theorem 1 also holds in the sense that a capacitated problem is at least as hard as the priority problem. In particular, given any priority cover instance $\text{PSC}(A, \pi, c)$ with k priorities, consider the capacitated instance $\text{CSC}(A, d, s, c)$ where each element j with priority p has demand $d_j = m^{2p}$ and a set i with priority p has supply m^{2p} , where m is the number of sets in A . It can be easily verified that a collection of sets is feasible for CSC if and only if it is feasible for PSC .

1.1 Our Results

Given a set system (X, S) with incidence matrix A , we will relate the integrality gap of a priority cover problem on A to the integrality gap of the set cover problem on A . We need the following additional definition.

Definition 1 (Hereditary Integrality Gap). A set system (X, S) with incidence matrix A has hereditary integrality gap α if the integrality gap for the natural LP relaxation of the set cover instance (A, c) restricted to any sub-system (X', S) , where $X' \subseteq X$, is at most α .

That is, the integrality gap is at most α if we restrict the system to any subset of elements.⁷ Clearly, solving separately for demands in each of the k priority classes, the integrality gap for any PSC instance is at most k times the hereditary integrality gap. We show that this can be improved substantially.

Theorem 2. *The integrality gap of any instance $\text{PSC}(A, \pi, c)$ with k priorities is $O(\alpha \log^2 k)$, where α is the hereditary integrality gap of the corresponding set cover instance (A, c) .*

According to Theorem 1, given any capacitated instance, the number of priorities in the associated priority instance is $k = O(\log s_{\max})$, and hence Theorem 2 implies that having capacities increases the hereditary integrality gap by at most $O((\log \log s_{\max})^2)$ (provided the multi-cover problem is also well-behaved w.r.t to the integrality gap).

Theorem 2 is proved in section 2 and its proof is surprisingly simple. However, this general result already achieves guarantees close to those known for very special systems. For example, for the previously mentioned CSC problem of covering points with rectangles touching the x -axis [1], a $O(\log k)$ guarantee for k priorities was obtained only recently using breakthrough geometric techniques of [13]. Using theorem 2 instead of the results of [13] already yields major improvements over previous results for the problem studied in [1].

Another corollary of Theorem 2 is that if A is *totally unimodular* (TU), then there this is an $O((\log \log s_{\max})^2)$ approximation for any capacitated problem on A . This follows as a TU matrix has a hereditary integrality gap of 1 for the multi-cover problem. This motivates our next result for set systems with low *hereditary discrepancy* (see Section 3 for a definition). Recall that TU matrices have a hereditary discrepancy of 1 (see e.g., [12]). Low hereditary discrepancy set systems arise naturally when the underlying system is a union of TU or other simpler systems. Recently, [8] also gave a surprising connection between low discrepancy and bin packing.

Theorem 3. *For any set system A where the dual set system A^T has hereditary discrepancy α , the integrality gap of the multi-cover instance $\text{MSC}(A, d, \mathbf{1}, c)$ for any demands d is α .*

As stated earlier, this implies an $O(\alpha(\log \log s_{\max})^2)$ integrality gap for any instance $\text{CSC}(A, d, s, c)$. Note that the integrality gap we show for the multi-cover problem is exactly α (and not just $O(\alpha)$), and hence this strictly generalizes the TU property, which results in an integral polytope.

The priority covering framework is particularly useful in geometric settings⁸. Appealing to this connection, our next result relates the VC dimension of the priority ver-

⁷ We note that this definition also allows the restriction of the system to $S' \subseteq S$, by considering the integrality gap on fractional solutions with support S' .

⁸ Given a geometric set cover problem, its priority version can be encoded as another geometric problem (this increases the underlying dimension by 1). By adding a new dimension to encode priority, replace sets S_i with priority p by the geometric object $S_i \times [0, p]$ and points \mathbf{p}_j with priority q by point $\mathbf{p}_j \times [0, q]$. It is easily checked that the set cover problem on this instance is equivalent to the priority cover problem on the original instance. This observation was used in [1, 5], and we do not elaborate more on it here.

sion of a problem to the original system. This is useful as low VC dimension can be exploited to obtain good set cover guarantees [3].⁹

Theorem 4. *For any set system with VC dimension d , the VC dimension of its priority version (for any setting of priorities) is at most $d + 1$.*

Lower Bounds. In light of the above results, two natural questions arise. First, can similar guarantees for capacitated version be obtained without any hereditary assumption, that is w.r.t to the integrality gap alone? Second, is the loss of factor $O(\log^2 k)$ in the guarantees necessary?

For the first question, we note that there are natural problems such as priority Steiner tree [6, 7], where the underlying set system is not hereditary, and the LP for the priority version has an integrality gap of $\Omega(k)$.

For the second question, in Section 5, we show that even for hereditary set systems, there are instances where the priority version has an integrality gap of $\Omega(\alpha \log k)$ when the original set cover problem has hereditary integrality gap of at most α .

Theorem 5. *There exist hereditary set cover instances with $O(1)$ -integrality gap for which the priority version has an integrality gap of $\Omega(\log k)$.*

These gap instances rely on the recent breakthrough constructions of Pach and Tardos [11] for geometric set systems with large ϵ -nets. In particular, we show that the gap already holds for the rectangle cover problem considered in [1] which we mentioned earlier. This shows that Theorem 2 is tight up to a $O(\log k)$ factor. Closing this gap would be an interesting question to study.

1.2 Other Related Work

Besides the work of Chakrabarty et al. [5] mentioned above, a work in spirit similar of ours is that of Kolliopoulos [10]. They studied the relationship between the approximability of a CSC and its corresponding set cover problem under the *no-bottleneck* assumption: this states that “the supply of every set/column is smaller than the demand of every element/row” (i.e. maximum supply is no more than minimum demand). Under this assumption, they show if the $x \leq 1$ constraint (or $x \leq d$ in general) can be violated by a constant multiplicative factor, then the integrality gap of any CSC is within an $O(1)$ factor of the corresponding $\{0, 1\}$ -CIP. However, nothing better than the standard set cover guarantee was known even with the no-bottleneck assumption. We refer the reader to [5] for further discussions on related work.

2 Bounding the Integrality gap of PSC’s

In this section, we prove Theorem 2. We show that the integrality gap of PSC instances that are characterized by hereditary set systems is $O(\alpha \log^2 k)$, where k is the number of priorities. Recall the stair-case structure of Π and the definition of $B = A \wedge \Pi$.

⁹ Note that we need to bound the VC dimension of the dual set system A^T to obtain guarantees for the set cover instance A .

The idea is rather simple. We decompose the incidence matrix B of the PSC instance into a collection of submatrices¹⁰ $\{D_0, \dots, D_\ell\}$ with the following properties.

1. Each such submatrix D_q is also a submatrix of A (and not just that of B).
2. Each element j appears as a row in at most $O(\log k)$ of the submatrices D_q and,
3. Each set i appears as a column in at most $O(\log k)$ of the submatrices D_q .

As A has a hereditary integrality gap of α , by the first property above, any fractional set cover solution restricted to the sub-system D_q has an integrality gap of α . Now, any fractional set cover solution x on A induces a fractional solution on D_q (after appropriate scaling). We use the second and the third properties stated above to show that the rounded solution for these D_q 's can be combined to obtain a feasible integral solution for A while increasing the cost by an $O(\log^2 k)$ factor. We begin by describing the decomposition procedure.

Decomposition Procedure: By adding dummy priorities if necessary, let us assume without loss of generality that k is an integral power of 2. Let the priorities be indexed by $1, \dots, k$ (with 1 being the lowest priority and k the highest). For priorities p, p', q, q' such that $p \leq p'$ and $q \leq q'$, let $B([p, p'][q, q'])$ denote the submatrix of B consisting of columns (resp. rows) with priorities in the range $[p, p']$ (resp. $[q, q']$).

A crucial observation is that if $p \geq q'$, then for any $p' \geq p$ and $q \leq q'$, the submatrix $B' = B([p, p'][q, q'])$ is also a submatrix of A . This simply follows as p is the lowest priority of any set in B' , which is at least as large as the priority of any row.

We define the decomposition of B inductively as follows: In the base case when $k = 1$, the decomposition consists of the single matrix $\{B\}$ itself. For general k , we define the decomposition as consisting of the matrix $D_0 = B([k/2 + 1, k][1, k/2])$, together with the (inductive) decompositions of

$$B_1 = B([1, k/2][1, k/2]) \quad \text{and} \quad B_2 = B([k/2 + 1, k][k/2 + 1, k]).$$

Note that both B_1 and B_2 involve only $k/2$ priorities. See Figure 2 for an illustration of the decomposition scheme.

Lemma 1. *The decomposition procedure satisfies the three properties claimed above.*

Proof. It is easily checked that this procedure gives a decomposition of B . Moreover, as $B([k/2 + 1, k][1, k/2])$ is a valid submatrix of A , it follows that all the submatrices obtained in the decomposition are submatrices of A .

Next we show by a simple induction that each element and set can lie in at most $1 + \log k$ submatrices D_q . This is clearly true if $k = 1$. Now, suppose $k > 1$ and consider some fixed element j . It can lie in the submatrix D_0 and exactly one of B_1 or B_2 . Since B_1 and B_2 are $k/2 \times k/2$ matrices, the claim follows by induction. An identical argument works for sets.

¹⁰ N is submatrix of M if N is obtained by restricting M to a subset of rows and columns.




Fig. 2. Recursive partitioning of B .

Rounding Algorithm:

1. Let $x^* = \{x_1^*, x_2^*, \dots, x_m^*\}$ be some optimal fractional solution for the set system B .
2. For each submatrix D in the decomposition of B , do the following:
 - (a) Let \mathcal{S}_D denote the collection of sets that lie in D , and let x_D denote the solution $\min(x^*(1 + \log k), 1)$ restricted to sets in \mathcal{S}_D .
 - (b) Let E_D be the set of elements in D that are covered fractionally to an extent of at least 1 by x_D .
 - (c) Consider the set system (E_D, \mathcal{S}_D) . Now, x_D is a feasible fraction set cover solution for this set system. As the hereditary integrality gap of A is α , apply the rounding algorithm to (E_D, \mathcal{S}_D) with x_D as the fractional solution. Let \mathcal{S}'_D denote the collection of sets chosen by this rounding.
3. Our final solution is simply the union of \mathcal{S}'_D , over all D in the decomposition of B .

Analysis: We first show that the algorithm produces a valid set cover and then bound the total cost, which will complete the proof of Theorem 2.

Lemma 2. *Each element in B is covered by some set in the solution.*

Proof. Consider some fixed element j . As j lies in at most $1 + \log k$ sets in the decomposition of B , and as x^* is a feasible fractional solution for B , there is some submatrix D that contains j and such that $\sum_{i \in \mathcal{S}_D} x_j^* A_{ij} \geq 1/(1 + \log k)$. Hence, the solution x_D covers j to an extent of at least 1 (i.e. $j \in E_D$), and the rounding algorithm applied to (E_D, \mathcal{S}_D) will ensure that j is covered by some set in \mathcal{S}'_D .

Lemma 3. *The total cost of the solution produced is $O(\log^2 k)\alpha$ times the LP cost.*

Proof. As A has hereditary integrality gap α , the cost of the collection \mathcal{S}'_D is at most α times the cost of the fraction solution x_D , which itself is at most $O(\log k)$ times the cost of solution x^* restricted to the variables (sets) in \mathcal{S}_D . As each set i lies in $O(\log k)$ submatrices D in the decomposition of B , summing up over all D , this implies that the total cost of the solution is $O(\alpha \log^2 k)$ times the cost of x^* .

3 Set Systems with small Hereditary Discrepancy

In this section, we consider set systems with low hereditary discrepancy and prove Theorem 3. Recall that, for a set system (X, \mathcal{S}) the discrepancy is defined as $\text{disc}(X, \mathcal{S}) = \min_f \max_{S' \in \mathcal{S}} |\sum_{e \in S'} f(e)|$ where $f : X \rightarrow \{-1, 1\}$ is a two coloring of the universe X , and the hereditary discrepancy is defined as $\text{herdisc}(X, \mathcal{S}) = \max_{X' \subseteq X} \text{disc}(X', \mathcal{S}|_{X'})$ where $\mathcal{S}|_{X'}$ is the collection of sets restricted to the elements X' .

In the setting of Theorem 3, where the (dual) set system A^T has *hereditary discrepancy* at most α , this means that given any sub-collection S' of sets, there is a $\{-1, +1\}$ coloring χ of S' that satisfies $|\sum_{i \in S'} A_{ij} \chi(i)| \leq \alpha$ for each row j .

3.1 Rounding Procedure

Let x^* be an optimal solution to the following natural LP relaxation of $\text{MSC}(A, d, \mathbf{1}, c)$.

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^m c_i x_i \\ & \text{subject to} && \sum_{i=1}^m A_{ij} x_i \geq d_j \quad \forall 1 \leq j \leq n \\ & && x_i \in [0, 1] \quad \forall S_i \in \mathcal{S} \end{aligned}$$

Scaling. First we scale x^* by a factor of α , i.e. $x'_i = \min(\alpha x_i^*, 1)$. Let \mathcal{H} be the set of variables for which $x' = 1$ and let $\mathcal{L} = \mathcal{S} \setminus \mathcal{H}$. Clearly, the solution $\{x'_i : i \in \mathcal{L}\}$ is feasible to the following (residual) set of constraints (for all elements j):

$$\sum_{i \in \mathcal{L}} A_{ij} x'_i \geq \alpha \left(d_j - \sum_{i \in \mathcal{H}} A_{ij} \right)$$

Iterative Rounding. In this step, we iteratively round the solution x' , without increasing its total cost, while also ensuring that the constraints remain satisfied. Consider the binary representation of variables in solution x' and let t denote the least significant bit in the representation. We index the rounds ℓ from t down to 1. Let us initialize the solution in the initial round $\ell = t$ as $x^t = x'$ and repeat the following step.

Round ℓ : Let S_ℓ denote the set of columns that have a 1 in their least significant bit (i.e. at position ℓ) in this round, and let $f_\ell : S_\ell \rightarrow \{-1, 1\}$ be a ± 1 coloring of the columns that minimizes discrepancy (w.r.t S_ℓ) for all the rows. Clearly, there exists one with discrepancy at most α .

Now, consider the following two solutions: For all $i \in S_\ell$, set $x_i^+ = x_i^\ell + \frac{f_\ell(i)}{2^\ell}$ and $x_i^- = x_i^\ell - \frac{f_\ell(i)}{2^\ell}$. As $x_i^+ + x_i^- = 2x_i^\ell$, it is easy to see that at least one of the solutions x^+ or x^- has cost no more than that of x^ℓ , and we set $x^{\ell-1}$ to that solution. Furthermore, because we have either added or subtracted $1/2^\ell$ from all the variables in S_ℓ , the least significant bit of the solution $x^{\ell-1}$ is now $\ell - 1$.

Having ensured that the cost does not increase, it remains to bound the change in the coverage of any element, for which we use the bounded discrepancy of the coloring f_ℓ . Indeed, since f_ℓ has discrepancy at most α , we have that $\sum_{i \in S_\ell} A_{ij} f_\ell(i) \in [-\alpha, \alpha]$

for all j , and hence

$$\sum_{i \in S_\ell} A_{ij} (x^{\ell-1} - x^\ell) = \sum_{i \in S_\ell} A_{ij} \frac{1}{2^\ell} f_\ell(i) \geq -\frac{\alpha}{2^\ell}.$$

Thus the coverage for any element drops by at most $\alpha/2^\ell$ in round ℓ , and this will be crucial for the analysis.

Output. By the invariant about the least significant bit after each round, at the conclusion of the rounding phase, all variables are either 0 or 1. Our final solution is then $\mathcal{H} \cup \mathcal{X}$, where $\mathcal{X} := \{i \in \mathcal{L} : x_i^0 = 1\}$.

3.2 Analysis

Final Cost. As the cost of the solution can only go down in each round ℓ , the cost of the final solution is at most that of x' , which is at most α times the LP optimum.

Feasibility. Consider any element j . In round ℓ , the coverage of j can drop by at most $\alpha/2^\ell$. Hence, over all the rounds, the total drop in coverage is $\sum_{i=1}^t \alpha/2^i$ which is strictly smaller than α . Therefore,

$$\sum_{i \in \mathcal{L}} A_{ij} x_i^0 > \alpha \left(d_j - \sum_{i \in \mathcal{H}} A_{ij} \right) - \alpha \geq \left(d_j - \sum_{i \in \mathcal{H}} A_{ij} - 1 \right). \quad (4)$$

As $\sum_{i \in \mathcal{L}} A_{ij} x_i^0$ is integral, the strict inequality in (4) implies that $\sum_{i \in \mathcal{L}} A_{ij} x_i^0 \geq (d_j - \sum_{i \in \mathcal{H}} A_{ij})$, and hence the solution is feasible.

4 Set Systems with small VC Dimension

We consider set systems with small VC dimension and prove theorem 4. We first recall the definition of VC dimension. Given a set system $(\mathcal{X}, \mathcal{S})$, for $X' \subseteq X$ let $\mathcal{S}_{|X'}$ denote the set system restricted to X' . We say that X' is shattered by \mathcal{S} if there are $2^{|X'|}$ distinct sets in $\mathcal{S}_{|X'}$. A set system (X, \mathcal{S}) is said to have VC dimension d , if $d > 0$ is the smallest integer such that no $d+1$ point subset $X' \subseteq X$ can be shattered. Also recall that the incidence matrix B of a PSC instance is obtained as $B_{ij} = A_{ij} \mathbf{1}_{\pi(S_i) \geq \pi(e_j)}$.

Theorem 4. *The VC dimension of the set system B is at most one more than that of A .*

Proof. Consider the matrix B and order the demand and the supply-priorities in non-decreasing order (as shown in Figure 1). Let d denote the VC dimension of A , and for the sake of contraction, suppose B has VC dimension at least $d+2$. Then there exists a subset of rows Y in B , $|Y| = d+2$, such that there are 2^{d+2} distinct columns in the submatrix induced by Y . Consider all the 2^{d+1} columns in this submatrix that have a 1 in their bottom-most coordinate. As $B_{ij} = A_{ij} \mathbf{1}_{\pi(S_i) \geq \pi(e_j)}$, every coordinate starting from the bottom-most coordinate with a 1 in B has the same value in both A and B . But then the rows of Y except the bottom-most one (there are $d+1$ of them) are shattered by A , contradicting that it has VC dimension d .

5 Lower Bounds

In this section, we establish a $\Omega(\log k)$ lower bound on the integrality gap of the PSC LP for hereditary instances for which the underlying set cover instance has $O(1)$ hereditary integrality gap. This shows that Theorem 2 is tight up to an $O(\log k)$ factor.

The Hinged Axis-Aligned Rectangle Cover Problem. The underlying problem we start with is the *Hinged Axis-Aligned Rectangle Cover* problem (HARC): we are given a set of points $X = \{(x_j, y_j) : 1 \leq j \leq n\}$ in the 2-dimensional plane and a collection of axis-aligned rectangles $\mathcal{S} = \{[a_i, b_i] \times [0, d_i] : 1 \leq i \leq m\}$ all of which have one side on the X-axis. The goal is to pick a minimum number of rectangles to cover X where the notion of coverage is simply containment of the point inside the rectangle.

It is known that the natural LP relaxation for this problem has an integrality gap of 2 [1]. Moreover the gap is clearly hereditary as any sub-collection of sets and elements is also a problem of the same type. We will now show that the natural LP relaxation for *Priority HARC* has an integrality gap of $\Omega(\log k)$ when there are k priorities. We achieve this by relating the priority version of the HARC problem to the 2D rectangle covering problem (2DRC).

The 2D Rectangle Cover Problem. In the 2DRC problem, we are given a set of points $X = \{(x_j, y_j) : 1 \leq j \leq n\}$ in the 2-dimensional plane and a collection of axis-aligned rectangles $\mathcal{S} = \{[a_i, b_i] \times [c_i, d_i] : 1 \leq i \leq m\}$. The goal is to pick a minimum number of rectangles to cover each of the given points where the notion of coverage is simply containment of the point inside the rectangle.

Step 1: Reducing 2DRC to Priority HARC. Consider an instance of 2DRC $\mathcal{I} = (X, \mathcal{S})$. Without loss of generality, we assume that no two points share any coordinate (which we can ensure by moving the points by infinitesimal amounts). We now create the Priority HARC instance \mathcal{I}' as follows: for each point $(x_j, y_j) \in X$, create the point $e_j = (x_j, y_j)$ with priority $\pi(e_j) = 1/y_j$. For each rectangle $[a_i, b_i] \times [c_i, d_i]$, create an axis-aligned rectangle $S_i = [a_i, b_i] \times [0, d_i]$ with priority $\pi(S_i) = 1/c_i$. By construction, it is clear to see that $(x_j, y_j) \in [a_i, b_i] \times [c_i, d_i]$ iff e_j is covered by S_i and $\pi(S_j) \geq \pi(e_i)$.

Note: Since each set could (in the worst case) be associated with its own priority, the number of priorities k created in the above reduction is $O(m)$, where m is the number of rectangles.

Step 2: Lower Bound for 2DRC. Therefore it suffices to obtain an integrality gap for the 2DRC in order to get the same gap for priority HARC. The idea is to use recent super-linear lower bounds on ϵ -nets for 2DRC, and the strong connection between ϵ -nets and the LP relaxation for set cover. In particular, we use the following theorem on ϵ -net lower bounds due to Pach and Tardos [11].

Theorem 6 ([11]). *For any $\epsilon > 0$ and for any sufficiently large integer $m \geq m_0(\epsilon) = \text{poly}(\frac{1}{\epsilon})$, there exists a range space (X, R) , where X is a set of points in \mathbb{R}^2 and R consists of m axis-aligned rectangles, such that the size of the smallest ϵ -net (w.r.t the points) is at least $\Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$. That is, if $S \subseteq R$ is such that any point $p \in X$ that is contained in at least ϵm rectangles is covered by a rectangle in S , then $|S| = \Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$.*

To get our integrality gap, consider the following set cover instance: the sets are all the rectangles, and elements are only those points which are contained in at least ϵm points. Then clearly from the above theorem, and the fact that any feasible integer solution is a valid ϵ -net, we get $\text{Opt} \geq \Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$. To complete the proof, we need to upper bound the cost of an optimal LP solution: if we set each x_S to $1/(\epsilon m)$, we see that such a solution is feasible, i.e., all elements are fractionally covered to extent 1; furthermore, the total cost of this fractional cover is $1/\epsilon$. Now this immediately gives us an integrality gap of $\Omega(\log(1/\epsilon))$. Now notice that the number of rectangles in the instances created can be set to $m = m_0(\epsilon) = \text{poly}(\frac{1}{\epsilon})$. Therefore, the integrality gap is also $\Omega(\log m) = \Omega(\log k)$, since m is linearly related to the number of priorities k as noted above. This proves Theorem 5.

Acknowledgments. We would like to thank Anupam Gupta for several useful discussions.

References

1. Nikhil Bansal and Kirk Pruhs. The geometry of scheduling. In *FOCS '10*, pages 407–414, 2010.
2. Amotz Bar-Noy, Reuven Bar-Yehuda, Ari Freund, Joseph (Seffi) Naor, and Baruch Schieber. A unified approach to approximating resource allocation and scheduling. *J. ACM*, 48:1069–1090, 2001.
3. H. Bronnimann and M. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete and Computational Geometry*, 14:463–479, 1995.
4. Robert D. Carr, Lisa K. Fleischer, Vitus J. Leung, and Cynthia A. Phillips. Strengthening integrality gaps for capacitated network design and covering problems. In *SODA '00*, pages 106–115, 2000.
5. Deeparnab Chakrabarty, Elyot Grant, and Jochen Konemann. On column-restricted and priority covering integer programs. In *IPCO '10*, pages 355–368, 2010.
6. Moses Charikar, Joseph (Seffi) Naor, and Baruch Schieber. Resource optimization in qos multicast routing of real-time multimedia. *IEEE/ACM Trans. Netw.*, 12:340–348, April 2004.
7. Julia Chuzhoy, Anupam Gupta, Joseph (Seffi) Naor, and Amitabh Sinha. On the approximability of some network design problems. In *SODA '05*, pages 943–951, 2005.
8. Friedrich Eisenbrand, Dömötör Pálvölgyi, and Thomas Rothvoß. Bin packing via discrepancy of permutations. In *SODA*, pages 476–481, 2011.
9. Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45:634–652, July 1998.
10. Stavros G. Kolliopoulos. Approximating covering integer programs with multiplicity constraints. *Discrete Appl. Math.*, 129:461–473, August 2003.
11. János Pach and Gábor Tardos. Tight lower bounds for the size of epsilon-nets. *CoRR*, abs/1012.1240, 2010.
12. A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
13. Kasturi Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In *STOC '10*, pages 641–648, 2010.