Tight Approximation for Partial Vertex Cover with Hard Capacities

Jia-Yau Shiau[1], Mong-Jen Kao[2], Ching-Chi Lin[3], and D.T. Lee[4]

1   **Graduate Institute of Electronics Engineering, National Taiwan University, Taipei, Taiwan**
    `janus7799@gmail.com`
2   **Institute of Information Science, Academia Sinica, Taipei, Taiwan**
    `mong@iis.sinica.edu.tw`
3   **Department of Computer Science and Engineering, National Taiwan Ocean University, Keelung City, Taiwan**
    `lincc@mail.ntou.edu.tw`
4   **Institute of Information Science, Academia Sinica, Taipei, Taiwan**
    `dtlee@ieee.org`

------- **Abstract** -------

We consider the partial vertex cover problem with hard capacity constraints (Partial VC-HC) on hypergraphs. In this problem we are given a hypergraph $G = (V, E)$ with a maximum edge size f and a covering requirement $\mathcal{R}$. Each edge is associated with a demand, and each vertex is associated with a capacity and an (integral) available multiplicity. The objective is to compute a minimum vertex multiset such that at least $\mathcal{R}$ units of demand from the edges are covered by the capacities of the vertices in the multiset and the multiplicity of each vertex does not exceed its available multiplicity.

In this paper we present an f-approximation for this problem, improving over a previous result of $(2f + 2)(1 + \epsilon)$ by Cheung et al to the tight extent possible. Our new ingredient of this work is a generalized analysis on the extreme points of the natural LP, developed from previous works, and a strengthened LP lower-bound obtained for the optimal solutions.

1   Introduction

The capacitated vertex cover problem with hard capacities (VC-HC) models a demand-to-service assignment scenario generalized from the classical vertex cover problem. In this problem, we are given a hypergraph $G = (V, E \subseteq 2^V)$ where each $e \in E$ is associated with a demand $d_e \in \mathbb{R}^{\geq 0}$ and each $v \in V$ is associated with a capacity $c_v \in \mathbb{R}^{\geq 0}$ and an available multiplicity $m_v \in \mathbb{Z}^{\geq 0}$. The objective is to find a vertex multiset, or, cover, represented by a demand assignment function $h \colon E \times V \to \mathbb{R}^{\geq 0}$, such that the followings are met:

(1) $\sum_{v \in e} h_{e,v} = 1$ for all $e \in E$,

(2) $x_v^{(h)} \leq m_v$ for all $v \in V$, where $x_v^{(h)} := \left\lceil \sum_{e \colon e \in E, \, v \in e} (d_e \cdot h_{e,v}) / c_v \right\rceil$,

and $\sum_{v \in V} x_v^{(h)}$ is minimized.

In this paper, we consider VC-HC with partial covering constraints (Partial VC-HC). Instead of dictating that all the demand be covered, we are specified a covering requirement $\mathcal{R}$ to be fulfilled. In particular, the constraint (1) above is replaced by the following two constraints:

$$(1.a) \quad \sum_{v \in e} h_{e,v} \leq 1 \text{ for all } e \in E, \qquad (1.b) \quad \sum_{e \in E} \sum_{v \in e} d_e \cdot h_{e,v} \geq \mathcal{R}.$$

Note that, under this notion, VC-HC is the special case for which we have $\mathcal{R} = \sum_{e \in E} d_e$. Moreover, following the convention used in the literature, it is equivalent to specify the slack of coverage $\mathcal{L}$, i.e., the amount of demand that is allowed to be left uncovered.

Background

The vertex cover problem is among the fundamental problems in the study of graph theory and approximation algorithms. It is known that an f-approximation can be obtained via LP rounding and duality, where f is the size of the largest hyperedge. Khot and Regev [13] showed that, assuming the unique game conjecture (UGC), approximating this problem to a ratio better than $f - \epsilon$ is NP-hard for any $\epsilon > 0$ and $f \geq 2$.

The capacitated vertex cover generalizes vertex cover in that demand-to-supply constraints are introduced in addition to the 0/1-covering model. Chuzhoy and Naor [4] considered VC-HC on simple graphs with unit edge demands, i.e., $f = 2$ and $d_e = 1$ for all $e \in E$. They presented a 3-approximation for this problem. They showed that, when the vertices are weighted, minimizing the weighted cost is at least set-cover-hard. Therefore $O(f)$-approximations towards weighted cost model is unlikely even for this simple setting.

Gandhi et al. [5] gave a 2-approximation for VC-HC with unit edge demand by a refined rounding approach to [4]. Saha and Khuller [15] considered general edge demands and presented an $O(f)$-approximation for f-hypergraphs. Cheung et al. [3] presented an improved approach for this problem. They presented a $(1 + 2/\sqrt{3})$-approximation for simple graphs and a $2f$-approximation for f-hypergraphs. The gap of approximation for VC-HC was closed recently by Kao and Wong [9,16] to a tight f-approximation for any $f \geq 2$.

The Partial VC-HC problem was first considered by Cheung et al in [3] and a $(2f+2)(1+\epsilon)$-approximation in time $O\left(|V|^{1/\epsilon}|E|\right)$ was presented, based on the rounding technique provided for VC-HC in the same paper and an exhaustive search on potential solutions with cardinality $O(1/\epsilon)$.

Related Work

When the available multiplicity of the vertices is unlimited, this problem is known as soft capacitated vertex cover (CVC). This problem was first considered by Guha et al. [7] and a 2-approximation was presented. Kao et al. [8,10,11] studied capacitated dominating set problem and presented a series of results for the complexity and approximability of this problem. Bar-Yehuda et al. [2] considered partial CVC and presented a 3-approximation for simple graphs. A tight approximation for Partial CVC was given by Mestre [14], based on a delicate primal-dual scheme developed for a gradually strengthened LP of this problem.

In additional to partial coverage, VC-HC with relaxed multiplicity constraints have also been addressed in the literature, i.e., the constraint $x_v \leq m_v$ for all v is relaxed in exchange of affordable solution quality. Grandoni et al. [6] considered VC-HC with weighted cost model and relaxed multiplicity constraints. For the case $m_v = 1$ for all $v \in V$, they showed that, when augmenting the available multiplicity by a factor of f, a cover with cost guarantee

of f^2 to the optimal feasible cover can be obtained.[1], The bi-criteria approximation ratio was recently improved by Kao et al. [12] to $\left(k, (1 + \frac{1}{k-1})(f-1)\right)$ for any $k \geq 2$ and arbitrary vertex multiplicities.

Our Contribution and Discussion.

In this paper we consider the partial VC-HC problem and present tight approximation result for this problem. Our main result is the following theorem.

▶ **Theorem 1.** *We can compute an f-approximation for partial VC-HC in polynomial time, where f is the size of the largest hyperedge.*

This improves over the previous ratio of $(2f+2)(1+\epsilon)$ in [3] for Partial VC-HC to the tight extent. Our algorithm builds upon the iterative rounding technique developed for the VC-HC problem in our previous works [9, 16]. In each iteration, our algorithm modifies the current instance based on the optimal extreme point solution of the current working LP. When certain structural property is attained for the gradually modified instance, the algorithm rounds up the solution and terminates.

In contrast to the previous works for VC-HC [9, 16], our new ingredient in this work comes in two-fold. The first one is a strengthened LP lower-bound for the optimal solution. Surprisingly, the natural LP formulation for Partial VC-HC, which performs arbitrarily badly even for very simple settings, can be tuned to give tight lower-bounds for the optimal solutions. This allows us to get rid of the exhaustive search step used in [3], which inevitably brings in an undesirable ϵ-dependency in their result.

Our second ingredient for this work is a new insight on the analysis of extreme point solutions of the extended natural LP. Instead of reducing instances of Partial VC-HC for f-hypergraphs into instances of VC-HC for $(f+1)$-hypergraphs and introducing an extra $O(1)$ constant in the approximation ratio, as was done in [3], we analyze the extreme point solution for the original instance directly. This keeps the possibility of tight approximation result alive. Together this gives our tight f-approximation for Partial VC-HC.

The rest of this paper is organized as follows. In Section 2 we formally define Partial VC-HC and the LP relaxation we will be using throughout this paper. In Section 3 we present our tight approximation algorithm for this problem. We conclude with an extension of our result and discussion for future direction in Section 4.

2 Problem Statement and LP Relaxation

In this section we formally define the Partial VC-HC problem and introduce the LP relaxation we will be using. Throughout this paper, we use $G = (V, E)$ to denote a hypergraph with vertex set V and edge set $E \subseteq 2^V$. Under this context, each hyperedge $e \in E$ is represented by the set of its incident vertices, which is a vertex subset of V. In other words, $e \subseteq V$ for all $e \in E$. We use $f := \max_{e \in E} |e|$ to denote the size of the largest hyperedge in the considered graph.

For any edge subset $\mathcal{E} \subseteq E$, we use $\mathcal{E}[v]$ to denote the set of edges in $\mathcal{E}$ that are incident to the vertex v. Formally, $\mathcal{E}[v] := \{ e : e \in \mathcal{E} \text{ such that } v \in e \}$.

---

[1] The bi-approximation ratio of [6] is updated due to the difference between the considered models.

2.1   Partial Vertex Cover with Hard Capacities (Partial VC-HC)

In this problem we are given a hypergraph $G = (V, E)$ and a covering requirement $\mathcal{R}$, where each $e \in E$ is associated with a demand d_e and each $v \in V$ is associated with a capacity c_v and an (integral) available multiplicities m_v.

A solution to this problem consists of a demand assignment function $h \colon E \times V \to \mathbb{R}^+ \cup \{0\}$, where $h_{e,v}$ denotes the fraction of the edge e that is assigned to the vertex v. The multiplicity of each vertex v given by h is denoted $x_v^{(h)} := \left\lceil \sum_{e \in E[v]} (d_e \cdot h_{e,v})/c_v \right\rceil$. The assignment h is *feasible* if   (1) $\sum_{v \in e} h_{e,v} \leq 1$ for all $e \in E$,   (2) $\sum_{e \in E} \sum_{v \in e} d_e \cdot h_{e,v} \geq \mathcal{R}$, and   (3) $x_v^{(h)} \leq m_v$ for all $v \in V$.

Given an instance $\Pi = (V, E, \mathcal{R}, \mathbf{c}, \mathbf{m}, \mathbf{d})$ as described above, the problem of Partial VC-HC is to find a feasible assignment h such that $\sum_{v \in V} x_v^{(h)}$ is minimized. Without loss of generality, we assume that the input graph G admits a feasible assignment since this condition can be checked via a max-flow computation.

For the ease of presentation, we also use $\mathcal{L} := \sum_{e \in E} d_e - \mathcal{R}$ to denote the amount of demand that can be left unassigned. Furthermore, for each $e \in E$, we use $h_e^\phi := 1 - \sum_{v \in e} h_{e,v}$ to denote the fraction that is left unassigned for edge e.

We remark that, when d_e and c_v are integer-valued for all $e \in E$, $v \in V$, by the integrality of b-matching polytope, any fractional assignment can be turned into an integral assignment, i.e., $d_e \cdot h_{e,v} \in \mathbb{Z}^{\geq 0}$ for all e, v, using a standard integer flow computation.

2.2   LP relaxation for Partial VC-HC

Given an instance $\Pi = (V, E, \mathbf{c}, \mathbf{m}, \mathbf{d})$ of Partial VC-HC, for

(1) a vertex subset $\mathcal{V} \subseteq V$,

(2) an edge subset $\mathcal{E} \subseteq E$,

(3) a residue fraction of the edges $\mathbf{r}$ to be covered, where $\mathbf{0} \leq \mathbf{r} \leq \mathbf{1}$, and

(4) an additional lower-bound $\boldsymbol{\ell}$, where $\mathbf{0} \leq \boldsymbol{\ell} \leq \mathbf{m}$, on the multiplicity of the vertices,

we consider the following relaxation, with $\Psi := (\mathcal{V}, \mathcal{E}, \mathbf{r}, \boldsymbol{\ell}, \mathbf{c}, \mathbf{m}, \mathbf{d})$ being a parameter tuple:

$$
\begin{aligned}
& \min_{(\mathbf{x}, \mathbf{h})} \quad \sum_{v \in \mathcal{V}} x_v && \text{LP}(\Psi) \\[1em]
& \text{s.t.} \\[0.5em]
& \quad \sum_{v \in e \cap \mathcal{V}} h_{e,v} + h_e^\phi = r_e, && \forall e \in \mathcal{E} && \text{(1a)} \\
& \quad \sum_{e \in \mathcal{E}[v]} d_e \cdot h_{e,v} \leq c_v \cdot x_v, && \forall v \in \mathcal{V} && \text{(1b)} \\
& \quad \sum_{e \in \mathcal{E}} d_e \cdot h_e^\phi \leq \mathcal{L}, && && \text{(1c)} \\
& \quad \ell_v \leq x_v \leq m_v, && \forall v \in \mathcal{V} && \text{(1d)} \\
& \quad 0 \leq h_{e,v} \leq x_v, && \forall e \in \mathcal{E}, \ v \in e \cap \mathcal{V} && \text{(1e)} \\
& \quad 0 \leq h_e^\phi, && \forall e \in \mathcal{E} && \text{(1f)}
\end{aligned}
$$

Since each of the variables $h_{e,v}$ and x_v is bounded from both below and above, the feasible region of $\text{LP}(\Psi)$ is a polytope, and the reference to its extreme points is well-defined.

Throughout this paper, for a given instance Π of Partial VC-HC, a number of different parameter tuples will be considered. However, since $\mathbf{m}$ and $\mathbf{d}$ will remain the same in every considered tuple, we will use $(\mathcal{V}, \mathcal{E}, \boldsymbol{r}, \boldsymbol{\ell}, \mathbf{c})$ to denote the parameter tuple Ψ for the LP.

2.3   The Integrality Gap

We have introduced an extended natural LP relaxation for Partial VC-HC in the previous section. It may seem that $\text{LP}(\Psi)$ with parameter tuple $\Psi = (V, E, \mathbf{1}, \mathbf{0}, \boldsymbol{c})$ has an unbounded integrality gap for Partial VC-HC.

Consider the following simple example, where we have one vertex v and one edge $e = \{v\}$. The capacity of v is k for some $k \in \mathbb{N}$, the demand of e is also k. The covering requirement $\mathcal{R}$ is 1. (Therefore, $\mathcal{L} = k - 1$.)

The optimal fractional solution to $\text{LP}(\Psi)$ has cost at most $1/k$ since we can set $x_v = h_{e,v} = 1/k$ and $h_e^\phi = (k-1)/k$. However, any integral solution has cost at least 1 since the vertex v has to be selected. Therefore the gap between the two solutions can be arbitrarily large, and it seems that we need an LP stronger than $\text{LP}(\Psi)$ to manage this problem.

In general, strengthening an LP for improved integrality gap can be a challenging task to accomplish. However, if we use the fact that any integral solution must also have an integral cost, then the lower-bound given by $\text{LP}(\Psi)$ for this simple example becomes $\lceil 1/k \rceil = 1$, which matches the integral solution. Although this may seem to be case-dependent, we will show in the rest of this paper that, surprisingly, this strengthened LP lower-bound is sufficient to give a tight f-approximation result for this problem.

3   Tight Approximation for Partial VC-HC

In this section we describe our tight approximation algorithm for Partial VC-HC. In each iteration, the algorithm makes local decisions based on current working LP and modifies the parameter tuple accordingly. When no such decisions are there to be made, it rounds up all vertices unconditionally and stops.

We first introduce notions and operations our algorithm will be using. In Section 3.1 we describe our algorithm in detail. The analysis is provided in Section 3.2.

Basic Notion and Operations

Let $\Psi = (\mathcal{V}, \mathcal{E}, \boldsymbol{r}, \boldsymbol{\ell}, \boldsymbol{c})$ be a parameter tuple and $p = (\boldsymbol{x}, \boldsymbol{h})$ be a feasible solution for $\text{LP}(\Psi)$.

▶ **Definition 2** (Supporting edge). For any $e \in \mathcal{E}$ and $v \in e$, we say that edge e *supports* vertex v in the solution p if $0 < h_{e,v} = x_v$.

The idea behind this definition is that, this condition gives information on how the rounding can be done locally. Suppose that $h_{e,v} = x_v$ holds for some e and v. Then from constraint (1b) of $\text{LP}(\Psi)$ we know that

$$d_e \cdot h_{e,v} \;\leq\; \sum_{e \in \mathcal{E}[v]} d_e \cdot h_{e,v} \;\leq\; c_v \cdot x_v$$

Since $h_{e,v} = x_v$, it follows that $d_e \leq c_v$. In other words, if we know that the vertex v is to be rounded up eventually, then its capacity will be sufficient for covering the demand of e. This suggests the edge-folding operation for our rounding algorithm.

▶ **Definition 3** (Edge folding). Let e be an edge that supports a vertex v in p. By *folding e into v*, we update Ψ as follows:

(i) Remove e from $\mathcal{E}$ and decrease c_v by d_e.

(ii) Impose the constraint $x_v \geq 1/f$ to LP(Ψ) by setting $\ell_v = 1/f$.

Notice that, from the discussion above, folding supporting edges into the supported vertices results in no loss in the feasibility and approximation guarantee of the final solution, provided that the supported vertices are to be rounded-up eventually.

▶ **Definition 4** (Vertex down-pinning). We say that a vertex $v \in \mathcal{V}$ is *tight* if $x_v = 1/f$. By pinning down the value of v, we update Ψ as follows:

(i) For each $e \in \mathcal{E}[v]$, decrease r_e by $h_{e,v}$.

(ii) Remove v from $\mathcal{V}$.

Intuitively a vertex is tight when its value hits the minimal extent that still results in no extra loss (in terms of an f-approximation) when rounded up. When a tight vertex is pinned down, we also pin down the assignment from its incident edges. By constraint (1e) it ensures that, for each of its incident edge, say, e, only a small amount of demand, i.e., $\leq 1/f$, is removed from r_e when v is removed from $\mathcal{V}$. This guarantees that e still has a relatively large assignment, i.e., $\geq 1/f$, to one of its remaining incident ends.

3.1 The Algorithm

In the following we describe our tight approximation algorithm for Partial VC-HC in detail. The algorithm begins with the initial parameter tuple $\Psi = (V, E, \mathbf{1}, \mathbf{0}, \mathbf{c})$. In each iteration, it proceeds in the four steps described below.

1. Let $\Psi = (\mathcal{V}, \mathcal{E}, \mathbf{r}, \boldsymbol{\ell}, \mathbf{c})$ denote the current parameter tuple.
   Solve LP(Ψ) for a *basic optimal solution* $p = (\mathbf{x}, \mathbf{h})$. Let $I := \left\{ v \in \mathcal{V} : 0 < x_v < \frac{1}{f} \right\}$.

2. If there exists an edge $e \in \mathcal{E}$ that supports some $v \notin I$ in p,
   then fold e into v.

3. If there exists a vertex $v \in \mathcal{V}$ that is tight,
   then pin down the value of v.

4. If any operation is performed in Step 2 or Step 3,
   then restart Step 1.
   Otherwise, round up vertices in V and stop.

Let $\hat{\Psi} = (\hat{\mathcal{V}}, \hat{\mathcal{E}}, \hat{\mathbf{r}}, \hat{\boldsymbol{\ell}}, \hat{\mathbf{c}})$ denote the parameter tuple when the algorithm enters the final rounding step and $\hat{p} = (\hat{x}, \hat{h})$ denote the basic optimal solution computed for $\hat{\Psi}$.

Let $\mathbf{h}'$ denote the $\{0, 1\}$-assignment function for $E \setminus \hat{\mathcal{E}}$ that indicates the vertex which each edge $e \in E \setminus \hat{\mathcal{E}}$ is folded into. In particular, for each $e \in E \setminus \hat{\mathcal{E}}$ and $v \in e$, the variable $h'_{e,v}$ is 1 if e is folded into v and 0 otherwise. Furthermore, for any $v \in V \setminus \hat{\mathcal{V}}$ and any $e \in \hat{\mathcal{E}}[v]$, let $h''_{e,v}$ denote the assignment value of e to v when v was pinned down.

The final output $(\mathbf{x}^*, \mathbf{h}^*)$ is defined as follows. For any $v \in V$ and $e \in E[v]$, let

$$x_v^* := \begin{cases} \lceil \hat{x}_v \rceil, & \text{if } v \in \hat{\mathcal{V}}, \\ 1, & \text{otherwise.} \end{cases} \quad \text{and} \quad h_{e,v}^* := \begin{cases} \hat{h}_{e,v}, & \text{if } e \in \hat{\mathcal{E}} \text{ and } v \in \hat{\mathcal{V}}, \\ h''_{e,v}, & \text{if } e \in \hat{\mathcal{E}} \text{ but } v \notin \hat{\mathcal{V}}, \\ h'_{e,v}, & \text{otherwise.} \end{cases}$$

Then $(\mathbf{x}^*, \mathbf{h}^*)$ is output as the approximate solution.

We remark that, since this approach is insensitive to multiple operations, in the actual algorithm we will fold every supporting edge and pin down the value of every tight vertex. Furthermore, ties are broken arbitrarily if an edge supports multiple vertices outside I.

Let Tight-Partial-VC-HC denote the above algorithm. Our tight approximation result is stated in the following theorem:

▶ **Theorem 5.** *On any instance* $\Pi = (V, E, \mathcal{R}, \boldsymbol{c}, \boldsymbol{m}, \boldsymbol{d})$ *of Partial VC-HC with maximum edge size* $f \geq 2$, *algorithm Tight-Partial-VC-HC computes an* f-*approximation* $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ *in polynomial time.*

3.2 Analysis

In this section we prove Theorem 5. First, since the algorithm iterates only if some edge is folded or some vertex is pinned down, we know that the number of iterations is at most $O(|E| + |V|)$. Let k denote the number of iterations the algorithm repeats before it enters the rounding stage. For $1 \leq i \leq k$, we use the following notations to denote the respective concepts we have in the i^{th} iteration:

- $\Psi^{(i)} = \left(\mathcal{V}^{(i)}, \mathcal{E}^{(i)}, \boldsymbol{r}^{(i)}, \boldsymbol{\ell}^{(i)}, \boldsymbol{c}^{(i)} \right)$: The parameter tuple the algorithm maintains when it enters the i^{th} iteration. Note that $\Psi^{(1)} = (V, E, \boldsymbol{1}, \boldsymbol{0}, \boldsymbol{c})$ is the initial tuple and $\Psi^{(k)} = \left(\hat{\mathcal{V}}, \hat{\mathcal{E}}, \hat{\boldsymbol{r}}, \hat{\boldsymbol{\ell}}, \hat{\boldsymbol{c}} \right)$ is the final tuple.
- $p^{(i)} = \left(\boldsymbol{x}^{(i)}, \boldsymbol{h}^{(i)} \right)$: The basic optimal solution computed for LP($\Psi^{(i)}$).

From the usage of edge-folding and vertex-pinning operations, it is not difficult to see that Algorithm Tight-Partial-VC-HC indeed produces a feasible solution for Partial VC-HC. We summarize the feasibility of the algorithm in the following lemma.

▶ **Lemma 6.** *Algorithm Tight-Partial-VC-HC outputs a feasible solution* $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ *for LP($\Psi^{(1)}$). Moreover,* $\boldsymbol{x}^*$ *is integral.*

To prove Lemma 6, it suffices to show that (1) After each iteration, the modified tuple is feasible for the next iteration, i.e., for any $1 \leq i < k$, the feasible region of LP($\Psi^{(i+1)}$) is not empty. Therefore the access to basic feasible solutions of the LP is always valid. (2) The solution (x^*, h^*) does not violate the constraints of LP($\Psi^{(1)}$).

For statement (1) above, the following lemma shows that, $\left(\boldsymbol{x}^{(i)}, \boldsymbol{h}^{(i)} \right)$ gives a feasible point for LP($\Psi^{(i+1)}$), thereby ensuring that its feasible region will not be empty.

▶ **Lemma 7.** *The solution* $\left(\boldsymbol{x}^{(i)}, \boldsymbol{h}^{(i)} \right)$, *when restricted to* $\mathcal{V}^{(i+1)}$ *and* $\mathcal{E}^{(i+1)}$, *is feasible for LP($\Psi^{(i+1)}$), for any* $1 \leq i < k$.

Lemma 7 and statement (2) are proved by verifying the corresponding LP constraints. In the rest of this section, we establish the approximation guarantee.

The Approximation Guarantee

From Lemma 7 we know that, the sum of fractional values over vertices in V between successive iterations, which is composed of the objective value of the LP and the values of the down-pinned vertices, will form a non-increasing sequence. Therefore the total fractional value we have in each iteration always gives a valid lower-bound to any optimal integral solution. In particular, we have the following lemma.

▶ **Lemma 8.** *For any $1 \le i \le k$,*

$$\left\lceil \frac{1}{f} \cdot \left|V \setminus \mathcal{V}^{(i)}\right| + \sum_{v \in \mathcal{V}^{(i)}} x_v^{(i)} \right\rceil \le OPT,$$

where OPT is the cost of any optimal integral solution.

Proof. Since OPT is integer-valued, it suffices to show that for any $1 \le i \le k$,

$$\frac{1}{f} \cdot \left|V \setminus \mathcal{V}^{(i)}\right| + \sum_{v \in \mathcal{V}^{(i)}} x_v^{(i)} \le OPT.$$

The base case $i = 1$ holds directly since $\mathcal{V}^{(1)} = V$ and since $\boldsymbol{x}^{(1)}$ is optimal for $\text{LP}(\Psi^{(1)})$.
For $1 < i \le k$, since $\mathcal{V}^{(i)} \subseteq \mathcal{V}^{(i-1)}$, by Lemma 7 and the optimality of $\boldsymbol{x}^{(i)}$ for $\text{LP}(\Psi^{(i)})$, we have

$$\sum_{v \in \mathcal{V}^{(i)}} x_v^{(i)} \le \sum_{v \in \mathcal{V}^{(i)}} x_v^{(i-1)}.$$

Therefore,

$$\frac{1}{f} \cdot \left|V \setminus \mathcal{V}^{(i)}\right| + \sum_{v \in \mathcal{V}^{(i)}} x_v^{(i)} \le \frac{1}{f} \cdot \left|V \setminus \mathcal{V}^{(i-1)}\right| + \left(\frac{1}{f} \cdot \left|\mathcal{V}^{(i-1)} \setminus \mathcal{V}^{(i)}\right| + \sum_{v \in \mathcal{V}^{(i)}} x_v^{(i-1)} \right)$$

$$= \frac{1}{f} \cdot \left|V \setminus \mathcal{V}^{(i-1)}\right| + \sum_{v \in \mathcal{V}^{(i-1)}} x_v^{(i-1)}.$$

Therefore this lemma follows by an induction on i. ◀

Note that, the rounding cost of vertices in $\left|V \setminus \mathcal{V}^{(k)}\right|$ is properly bounded by its own fractional value since the fractional value of each down-pinned vertex is exactly $1/f$. Therefore it suffices to show that the rounding cost of vertices in $\mathcal{V}^{(k)}$ can be properly bounded as well.

Let $I^{(k)} := \left\{ v \in \mathcal{V}^{(k)} : 0 < x_v^{(k)} < \frac{1}{f} \right\}$ and $M^{(k)} := \left\{ v \in \mathcal{V}^{(k)} : 0 < x_v^{(k)} = m_v \right\}$ denote

the set of vertices with small and large fractional values in $p^{(k)}$, respectively. The following lemma upper-bounds the cardinality of $I^{(k)}$ in terms of the cardinality of $M^{(k)}$.

▶ **Lemma 9.** $\left|I^{(k)}\right| \le \left|M^{(k)}\right| + 1.$

Proof. Recall that $p^{(k)} = (\boldsymbol{x}^{(k)}, \boldsymbol{h}^{(k)})$ is an extreme point solution of $\text{LP}(\Psi^{(k)})$. Therefore, it follows that the rank of the tight constraints we have, i.e., the maximum number of linearly independent constraints we can pick among the inequalities that hold with equality, w.r.t. $p^{(k)}$ equals the total number of variables.

Since there exist no tight vertices in $\mathcal{V}^{(k)}$ nor edges in $\mathcal{E}^{(k)}$ that supports vertices outside $I^{(k)}$ by the design of the algorithm, it follows that, each tight constraint we have in $\text{LP}(\Psi^{(k)})$ w.r.t. $p^{(k)}$ must belong to one of the forms listed in Figure 1.

Modify the tight constraints as follows: Remove all zero-valued varaibles from the constraints of types (2a), (2b), and (2c). Plug the constraints of types (2d) and (2e) back into that of (2a) and (2b) by replacing corresponding x_v with m_v and $h_{e,v}$ with x_v.

Note that, by doing so, we literally remove equal number of variables and linearly independent tight constraints from the above. Therefore, the number of remaining variable still equals the rank of remaining tight constraints.

$$\sum_{v \in e \cap \mathcal{V}^{(k)}} h_{e,v} + h_e^{\phi} = r_e^{(k)}, \qquad e \in \mathcal{E}^{(k)} \tag{2a}$$

$$\sum_{e \in \mathcal{E}^{(k)}[v]} d_e \cdot h_{e,v} = c_v^{(k)} \cdot x_v, \qquad v \in \mathcal{V}^{(k)} \tag{2b}$$

$$\sum_{e \in \mathcal{E}^{(k)}} d_e \cdot h_e^{\phi} = \mathcal{L}, \tag{2c}$$

$$x_v = m_v, \qquad v \in M^{(k)} \tag{2d}$$

$$h_{e,v} = x_v, \qquad e \in \mathcal{E}^{(k)}, \ v \in e \cap I^{(k)} \tag{2e}$$

$$x_v = 0, \qquad v \in \mathcal{V}^{(k)} \tag{2f}$$

$$h_{e,v} = 0, \qquad e \in \mathcal{E}^{(k)}, \ v \in e \cap \mathcal{V}^{(k)} \tag{2g}$$

$$h_e^{\phi} = 0, \qquad e \in \mathcal{E}^{(k)} \tag{2h}$$

■ **Figure 1** Possible tight constraints we have in $\mathrm{LP}(\Psi^{(k)})$ with reference to the solution $p^{(k)}$.

Define the following sets:

- $H^* := \big\{ (e,v) : e \in \mathcal{E}^{(k)}, \ v \in e \cap I^{(k)}, \ 0 < h_{e,v} = x_v \big\}$. This corresponds to the non-zero $h_{e,v}$ variables that have been replaced by x_v due to constraints of type (2e).
- $H^+ := \big\{ (e,v) : e \in \mathcal{E}^{(k)}, \ v \in e \cap \mathcal{V}^{(k)}, \ 0 < h_{e,v} < x_v \big\}$, which corresponds to the set of $h_{e,v}$ variables that remains in the constraints.
- $\mathcal{V}^+ := \big\{ v : v \in \mathcal{V}^{(k)} \setminus M^{(k)}, \ x_v > 0 \big\}$ and $E^{\phi} := \big\{ e : e \in \mathcal{E}^{(k)}, \ h_e^{\phi} > 0 \big\}$.

The remaining tight constraints we have, after modification, are of the following forms, where $\tilde{c}_v^{(k)} = c_v^{(k)} - \sum_{e:\, (e,v) \in H^*} d_e$ is the resulting constant after the constraints (2e) are plugged in.

$$\sum_{v:\, (e,v) \in H^+} h_{e,v} + \sum_{v:\, (e,v) \in H^*} x_v + \big[e \in E^{\phi} \big] \cdot h_e^{\phi} = r_e^{(k)}, \qquad e \in \mathcal{E}^{(k)} \tag{3a}$$

$$\sum_{e:\, (e,v) \in H^+} d_e \cdot h_{e,v} = \tilde{c}_v^{(k)} \cdot x_v, \qquad v \in \mathcal{V}^+ \tag{3b}$$

$$\sum_{e \in E^{\phi}} d_e \cdot h_e^{\phi} = \mathcal{L}, \tag{3c}$$

$$\sum_{e:\, (e,v) \in H^+} d_e \cdot h_{e,v} = c_v^{(k)} \cdot m_v, \qquad v \in M^{(k)} \tag{3d}$$

In the following we work with constraints (3a) through (3d). We show that, the cardinality of $I^{(k)}$ is at most the number of constraints of types (3c) and (3d).

Let $\mathcal{C}$ be a maximal collection of linearly independent constraints from (3a) through (3d) and Vars denote the set of remaining variables. It follows that $|\mathcal{C}| = |\mathrm{Vars}| = |H^+| + |\mathcal{V}^+| + |E^{\phi}|$ and the coefficient matrix of $\mathcal{C}$ is of full-rank. Therefore, by identifying a set of pivots

of the coefficient matrix, we obtain a bijection $\sigma\colon \text{Vars} \mapsto \mathcal{C}$ such that any variable, say, $a \in \text{Var}$, appears in the constraint $\sigma(a)$.

For the brevity of notations, in the following for any $e \in \mathcal{E}^{(k)}$ we will use e to denote the corresponding type (3a) constraint it represents. Similarly, for any $v \in \mathcal{V}^{(k)}$, we use v to denote the type (3b) or (3d) constraint it represents. We use ϕ to denote the only one type (3c) constraint.

We have the following two properties, obtained from the fact that σ is a bijection.

▶ **Claim 10.** *We have $\sigma(x_v) = v$ for any $v \in \mathcal{V}^+ \setminus I^{(k)}$.*

Proof. Consider any $v \in \mathcal{V}^+ \setminus I^{(k)}$. Since $(e, v) \notin H^*$ for all e, the only constraint that contains x_v is v itself. Therefore σ must map x_v to v.   ◀

▶ **Claim 11.** *For any $e \in \mathcal{E}^{(k)}$, either $h_e^\phi > 0$ or $h_{e,v} > 0$ for some $v \in \mathcal{V}^{(k)} \setminus I^{(k)}$.*

Proof. Suppose that $h_e^\phi = 0$. We will argue that $h_{e,v} > 0$ for some $v \in \mathcal{V}^{(k)} \setminus I^{(k)}$.

Since the value of each down-pinned vertex equals $1/f$, the residue fraction of e decreases by at most $1/f$ each time when one of its incident vertices is pinned down. Therefore it follows that

$$r_e^{(k)} \ \geq \ 1 - \frac{1}{f} \cdot \left| e \cap \left(V \setminus \mathcal{V}^{(k)} \right) \right| \ \geq \ \frac{1}{f} \cdot \left| e \cap \mathcal{V}^{(k)} \right|,$$

where the last inequality follows from the fact that $|e| \leq f$.

This says,

$$\sum_{v\colon (e,v) \in H^+} h_{e,v} \ + \sum_{v\colon (e,v) \in H^*} x_v \ \geq \ \frac{1}{f} \cdot \left| e \cap \mathcal{V}^{(k)} \right|.$$

Therefore, one of the variables in the L.H.S. has value at least $1/f$, and it must be $h_{e,v'}$ for some v' with $(e, v') \in H^+$ since $x_v < 1/f$ for all v with $(e, v) \in H^*$. Since $h_{e,v'} \geq 1/f$, we know that $v' \notin I^{(k)}$ and this claim follows.   ◀

Consider any $v \in I^{(k)}$ and the constraint $\sigma(x_v)$. Since the variable x_v appears only in (3a) and (3b), we have the following two cases.

- If $\sigma(x_v) = e$ for some $e \in \mathcal{E}^{(k)}$, then according to Claim 11, either $h_e^\phi > 0$ or there exists $v' \in \mathcal{V}^{(k)} \setminus I^{(k)}$ such that the variable $h_{e,v'}$ appears in e.
  If $h_e^\phi > 0$, then σ must map h_e^ϕ to ϕ, i.e., $\sigma(h_e^\phi) = \phi$, since the variable h_e^ϕ appears only in e and ϕ but e is already mapped to by x_v.
  Similarly, if $h_{e,v'} > 0$ for some $v' \in \mathcal{V}^{(k)} \setminus I^{(k)}$, then it follows that $\sigma(h_{e,v'}) = v'$. Therefore, by Claim 10, v' does not belong to $\mathcal{V}^+ \setminus I^{(k)}$ and v' must belong to $I^{(k)} \cup M^{(k)}$. Hence we know that $v' \in M^{(k)}$.
  Therefore, if $\sigma(x_v) = e$ for some $e \in \mathcal{E}^{(k)}$, then x_v corresponds exclusively to either ϕ or a vertex $v' \in M^{(k)}$.
- If $\sigma(x_v) = v$, then there exists some e such that $(e, v) \in H^+$ since the constraint v is linearly independent from the others, implying that constraint v is non-degenerating. Since $h_{e,v}$ appears only in constraints e and v, it follows that $\sigma(h_{e,v}) = e$. By the same argument as the previous case, we can again associate x_v exclusively with either ϕ or some vertex $v' \in M^{(k)}$.

From the two cases above, we can associate each $v \in I^{(k)}$ exclusively to either ϕ or some vertex in $M^{(k)}$. It follows that $\left| I^{(k)} \right| \leq \left| M^{(k)} \right| + 1$ and this lemma is proved.   ◀

The following lemma bounds the cost of $\boldsymbol{x}^*$, which is exactly $\left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in \mathcal{V}^{(k)}} \left\lceil x_v^{(k)} \right\rceil$, and establishes the approximation guarantee.

▶ **Lemma 12.** *We have*

$$\left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in \mathcal{V}^{(k)}} \left\lceil x_v^{(k)} \right\rceil \leq f \cdot OPT.$$

Proof. By Lemma 8 it suffices to show that

$$\left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in \mathcal{V}^{(k)}} \left\lceil x_v^{(k)} \right\rceil \leq f \cdot \left\lceil \frac{1}{f} \cdot \left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in \mathcal{V}^{(k)}} x_v^{(k)} \right\rceil. \tag{4}$$

Note that, since $\left\lceil x_v^{(k)} \right\rceil \leq f \cdot x_v^{(k)}$ for any $v \in \mathcal{V}^{(k)} \setminus I^{(k)}$, this inequality holds trivially if $I^{(k)} = \emptyset$. In the following we assume that $I^{(k)} \neq \emptyset$.

Pair up the vertices in $I^{(k)}$ with that in $M^{(k)}$ by associating each $v \in I^{(k)}$ a distinct vertex from $M^{(k)}$. By Lemma 9, at most one vertex in $I^{(k)}$ could be unpaired. Without loss of generality, we assume that exactly one vertex, say, u, from $I^{(k)}$ is unpaired. It follows that

$$\sum_{v \in I^{(k)} \cup M^{(k)}} \left\lceil x_v^{(k)} \right\rceil \leq \left\lceil x_u^{(k)} \right\rceil + \sum_{v \in M^{(k)}} (m_v + 1) \leq \left\lceil x_u^{(k)} \right\rceil + f \cdot \sum_{v \in M^{(k)}} m_v, \tag{5}$$

where the last inequality holds since $f \geq 2$ and since $m_v \geq 1$ for all $v \in M^{(k)}$.

Let $W := \mathcal{V}^{(k)} \setminus \left(I^{(k)} \cup M^{(k)}\right)$. Plugging Inequality (5) into the L.H.S. of Inequality (4), since $\sum_{v \in M^{(k)}} m_v$ is integral and can be moved freely outside the ceilings, we can conclude that, Inequality (4) holds if the following statement holds.

$$\frac{1}{f} \cdot \left(\left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in W} \left\lceil x_v^{(k)} \right\rceil + \left\lceil x_u^{(k)} \right\rceil \right) \leq \left\lceil \frac{1}{f} \cdot \left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in W} x_v^{(k)} + x_u^{(k)} \right\rceil. \tag{6}$$

In the following we establish Inequality (6). Let p and q denote the integer multiple and the remainder of $\left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in W} \left\lceil x_v^{(k)} \right\rceil$ with the modulus f, i.e., we write

$$\left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in W} \left\lceil x_v^{(k)} \right\rceil = p \cdot f + q,$$

where $p, q \in \mathbb{Z}^{\geq 0}$, $0 \leq q < f$. The L.H.S. of Inequality (6) simplifies to

$$\frac{1}{f} \cdot \left(pf + q + 1 \right) \leq p + 1 = \left\lceil \frac{1}{f} \cdot \left(\left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in W} \left\lceil x_v^{(k)} \right\rceil \right) + x_u^{(k)} \right\rceil$$

$$\leq \left\lceil \frac{1}{f} \cdot \left|V \setminus \mathcal{V}^{(k)}\right| + \sum_{v \in W} x_v^{(k)} + x_u^{(k)} \right\rceil,$$

where the second equality holds since $0 < x_u^{(k)} < 1/f$ and the last inequality follows from the fact that $1/f < x_v^{(k)} < m_v$ for each $v \in W$, which implies that $(1/f) \cdot \left\lceil x_v^{(k)} \right\rceil \leq x_v^{(k)}$. This proves Inequality (6) and this lemma follows. ◀

4 Conclusion

We conclude with an interesting generalization of Partial VC-HC where we are given, instead of one global coverage requirement, a set of partial coverage constraints to be satisfied. In particular, we are given a function $R\colon 2^E \to [0,1]$, where for any $A \subseteq E$, $R(A)$ denotes the fraction of demand coverage among the edges in A to be fulfilled.[2]

Let $k = \big| \big\{ A : A \subseteq E, R(A) > 0 \big\} \big|$ denote the number of partial coverage constraints. By a direct generalization of Lemma 9 combined with the exhaustive search technique used in [3], we can obtain an $(f + k)(1 + \epsilon)$-approximation in $poly\left(|V|^{k/\epsilon}|E|\right)$ time.

However, it is not yet clear how better guarantees can be obtained nor how the dependency of k could be removed, and this would be an interesting direction to explore. The MFN-type LP relaxations, originally developed for capacitated facility location [1], could be a powerful tool to manage the approximation guarantee.

── **References** ─────────────────────────────

1   Hyung-Chan An, Mohit Singh, and Ola Svensson. Lp-based algorithms for capacitated facility location. *SIAM J. Comput.*, 46(1):272–306, 2017. `doi:10.1137/151002320`.

2   R. Bar-Yehuda, G. Flysher, J. Mestre, and D. Rawitz. Approximation of partial capacitated vertex cover. *SIAM Journal on Discrete Mathematics*, 24(4):1441–1469, 2010. `doi:10.1137/080728044`.

3   W.-C. Cheung, M. Goemans, and S. Wong. Improved algorithms for vertex cover with hard capacities on multigraphs and hypergraphs. In *Proceedings of SODA'14*, 2014. `doi:10.1137/1.9781611973402.124`.

4   Julia Chuzhoy and Joseph Naor. Covering problems with hard capacities. *SIAM Journal on Computing*, 36(2):498–515, August 2006. `doi:10.1137/S0097539703422479`.

5   R. Gandhi, E. Halperin, S. Khuller, G. Kortsarz, and A. Srinivasan. An improved approx algorithm for vertex cover with hard capacities. *J. Comput. Syst. Sci.*, 72:16–33, 2006. `doi:10.1016/j.jcss.2005.06.004`.

6   F. Grandoni, J. Könemann, A. Panconesi, and M. Sozio. A primal-dual bicriteria distributed algorithm for capacitated vertex cover. *SIAM J. Comput.*, 38(3), 2008. `doi:10.1137/06065310X`.

7   Sudipto Guha, Refael Hassin, Samir Khuller, and Einat Or. Capacitated vertex covering. *Journal of Algorithms*, 48(1):257–270, August 2003. `doi:10.1016/S0196-6774(03)00053-1`.

8   Mong-Jen Kao. *An Algorithmic Approach to Local and Global Resource Allocations*. PhD thesis, National Taiwan University, 2012.

9   Mong-Jen Kao. Iterative partial rounding for vertex cover with hard capacities. In *Proceedings of SODA'17*, pages 2638–2653, 2017. URL: `http://dl.acm.org/citation.cfm?id=3039686.3039860`.

10   Mong-Jen Kao, Han-Lin Chen, and D.T. Lee. Capacitated domination: Problem complexity and approximation algorithms. *Algorithmica*, November 2013. `doi:10.1007/s00453-013-9844-6`.

11   Mong-Jen Kao, Chung-Shou Liao, and D. T. Lee. Capacitated domination problem. *Algorithmica*, 60(2):274–300, June 2011. `doi:10.1007/s00453-009-9336-x`.

─────────────

[2]   Under this notion, Partial VC-HC is the special case for which we have $R(E) := \mathcal{R} / \sum_{e \in E} d_e$ and $R(A) := 0$ for any $A \subset E$.

12   Mong-Jen Kao, Hai-Lun Tu, and D. T. Lee. O(f) bi-approximation for capacitated covering
        with hard capacities. In *ISAAC'16*, pages 40:1–40:12, 2016. `doi:10.4230/LIPIcs.ISAAC.`
        `2016.40`.

13   Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within
        $2 - \epsilon$. *Journal of Computer and System Sciences*, 74(3):335–349, May 2008. `doi:10.1016/`
        `j.jcss.2007.06.019`.

14   Julián Mestre. A primal-dual approximation algorithm for partial vertex cover: Making
        educated guesses. *Algorithmica*, 55(1):227–239, 2009. `doi:10.1007/s00453-007-9003-z`.

15   B. Saha and S. Khuller. Set cover revisited: Hypergraph cover with hard capacities. In
        Proceedings of ICALP'12, pages 762–773, 2012. `doi:10.1007/978-3-642-31594-7_64`.

16   Sam Chiu-wai Wong. Tight algorithms for vertex cover with hard capacities on multigraphs
        and hypergraphs. In *Proceedings of SODA'17*, pages 2626–2637, 2017. URL: `http://dl.`
        `acm.org/citation.cfm?id=3039686.3039859`.