

$O(f)$ Bi-Approximation for Capacitated Covering with Hard Capacities*

Mong-Jen Kao¹, Hai-Lun Tu², and D. T. Lee³

- 1 Institute of Information Science, Academia Sinica, Taipei, Taiwan
- 2 Department of Computer Science and Information Engineering, National Taiwan University, Taipei, Taiwan
- 3 Institute of Information Science, Academia Sinica, Taipei, Taiwan; and Department of Computer Science and Information Engineering, National Taiwan University, Taipei, Taiwan

Abstract

We consider capacitated vertex cover with hard capacity constraints (VC-HC) on hypergraphs. In this problem we are given a hypergraph $G = (V, E)$ with a maximum edge size f . Each edge is associated with a demand and each vertex is associated with a weight (cost), a capacity, and an available multiplicity. The objective is to find a minimum-weight vertex multiset such that the demands of the edges can be covered by the capacities of the vertices and the multiplicity of each vertex does not exceed its available multiplicity.

In this paper we present an $O(f)$ bi-approximation for VC-HC that gives a trade-off on the number of augmented multiplicity and the cost of the resulting cover. In particular, we show that, by augmenting the available multiplicity by a factor of $k \geq 2$, a cover with a cost ratio of $(1 + \frac{1}{k-1})(f-1)$ to the optimal cover for the original instance can be obtained. This improves over a previous result, which has a cost ratio of f^2 via augmenting the available multiplicity by a factor of f .

1998 ACM Subject Classification G.2.2 Graph Algorithms

Keywords and phrases Capacitated Covering, Hard Capacities, Bi-criteria Approximation

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2016.40

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 $G = (V, E \subseteq 2^V)$ with maximum edge size f , where each $e \in E$ satisfies $|e| \leq f$ and is associated with a demand $d_e \in \mathbb{R}^{\geq 0}$, and each $v \in V$ is associated with a weight (or cost) $w_v \in \mathbb{R}^{\geq 0}$, 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 $h: E \times V \rightarrow \mathbb{R}^{\geq 0}$, such that the following two constraints are met:

1. $\sum_{v \in e} h_{e,v} \geq d_e$ for all $e \in E$,
 2. $x_v^{(h)} \leq m(v)$ for all $v \in V$, where $x_v^{(h)} := \left\lceil \sum_{e: e \in E, v \in e} h_{e,v} / c_v \right\rceil$,
- and $\sum_{v \in V} w(v) \cdot x_h(v)$ is minimized.

* This work is supported in part by Ministry of Science and Technology (MOST), Taiwan, under Grant MOST104-2218-E-001-032.



In this paper, we consider bicriteria approximation for VC-HC with augmented multiplicity constraints. In particular, we say that a demand assignment h forms an augmented (β, γ) -cover if it is feasible for the augmented multiplicity function $m'_v := \beta \cdot m_v$ for all $v \in V$ and the cost ratio is at most γ compared to the optimal assignment for the original instance. In other words, we are allowed to use additional multiplicities of the vertices up to a factor of β .

Background and Prior Work

The capacitated vertex cover generalizes vertex cover in that a demand-to-service assignment model is evolved from the original 0/1 covering model. This transition was exhibited via several work.

For classical vertex cover, it is known that a f -approximation can be obtained by LP rounding and duality [1, 8]. Khot and Regev [13] showed that, assuming the unique game conjecture, approximating this problem to a ratio better than $f - \epsilon$ is NP-hard for any $\epsilon > 0$ and $f \geq 2$.

Chuzhoy and Naor [4] considered VC-HC on simple graphs with unit edge demands, i.e., $|e| = 2$ and $d_e = 1$ for all $e \in E$. They presented a 3-approximation for the unweighted version of this problem, i.e., $w_v = 1$ for all $v \in V$. On the contrary, they showed that the weighted version is at least as hard as set cover, which renders $O(f)$ -approximations unlikely to exist even for this simple setting. Due to this reason, subsequent work on VC-HC has focused primarily on the unweighted version.

Gandhi et al. [5] gave a 2-approximation for unweighted VC-HC with unit edge demand by presenting a refined rounding approach to [4]. Saha and Khuller [14] 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 this problem was recently closed by Kao [10], who presented an f -approximation for any $f \geq 2$.

Grandoni et al. [6] considered weighted VC-HC with unit vertex multiplicity, i.e., $m_v = 1$ for all $v \in V$, and augmented multiplicity constraints. They presented a primal-dual approach that yields an augmented $(2, 4)$ -cover for simple graphs¹, which further extends to augmented (f, f^2) -cover for f -hypergraphs. This approach does not generalize, however, to arbitrary vertex multiplicities and does not entail further parametric trade-off either.

Further Related Work

The capacitated covering problem has been studied in various forms and variations. When the number of available multiplicities is unlimited, this problem is referred to soft capacitated vertex cover (CVC). This problem was first considered by Guha et al. [7], who gave a 2-approximation based on primal-dual. Kao et al. [11, 12, 9] 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 based on local ratio techniques.

Wolsey [15] considered submodular set cover, which includes classical set cover as a special case and which relates to capacitated covering in a simplified form, and presented a $(\ln \max_S f(S) + 1)$ -approximation. This approach was generalized by Chuzhoy and Naor [4]

¹ The bicriteria approximation ratio of [6] is updated in the context due to the different considered models. In [6] each vertex is counted at most once in the cost of the cover, disregarding the number of multiplicities it needs. In our model, however, the cost is weighted over the multiplicities of each vertex.

to capacitated set cover with hard capacities and unit demands, for which a $(\ln \delta + 1)$ -approximation was presented, where δ is the maximum size of the sets.

Our Result and Approach

We consider VC-HC with general parameters and present bicriteria approximations that yields a trade-off between the number of augmented multiplicities and the resulting cost. Our main result is the following bicriteria approximation algorithm:

► **Theorem 1.** *For any integer $k \geq 2$, we can compute an augmented $\left(k, \left(1 + \frac{1}{k-1}\right)(f-1)\right)$ -cover for weighted VC-HC in polynomial time.*

This improves over the previous ratio of (f, f^2) in [6] and provides a parameter trade-off on the augmented multiplicity and the quality of the solution. In particular, the cost ratio we obtained for this bi-approximation is bounded within $\frac{3}{2}(f-1)$ for all $k \geq 2$ and converges asymptotically to $f-1$ as k tends to infinity.

Our algorithm builds on primal-dual charging techniques combined with a flow-based procedure that exploits the duality of the LP relaxation. The primal-dual scheme we present extends the basic framework from [12, 7], which were designed for the soft capacity model where $m_v = \infty$ for all v . In contrast to the previous result in [6], we employ a different way of handling the dual variables as well as the primal demand assignments that follow. The seemingly subtle difference entails dissimilar analysis and gives a guarantee that is unavailable via their approach.

In particular, for the primal demand assignments, we use flow-based arguments to deal with pending decisions. This ensures that the vertices whose multiplicity limits are attained receive sufficient amount of demands to pay for their costs. The crucial observation in establishing the bicriteria approximation factor is that the feasible regions of the dual LP remains unchanged when the multiplicity constraint is augmented. Therefore the cost of the solution obtained via the primal-dual approach can be bounded by the optimal cost of the original instance. Together this gives our bi-approximation result.

The rest of this paper is organized as follows. In §2 we formally define VC-HC and introduce the natural LP relaxation and its dual LP for which we will be working with. For a better flow to present our bicriteria approximation, we first introduce our primal-dual algorithm and the corresponding analysis in §3. In §4 we establish the bi-approximation approximation ratio and prove Theorem 1. Finally we conclude in §5 with some future directions for related problems.

2 Problem Statement and LP Relaxation

Let $G = (V, E)$ denote a hypergraph with vertex set V and edge set $E \subseteq 2^V$ and $f := \max_{e \in E} |e|$ denote the size of the largest hyperedge in G . For any $v \in V$, we use $E[v]$ to denote the set of edges that are incident to the vertex v . Formally, $E[v] := \{e : e \in E \text{ such that } v \in e\}$. This definition extends to set of vertices, i.e., for any $A \subseteq V$, i.e., $E[A] := \bigcup_{v \in A} E[v]$.

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

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 weight (or cost) $w_v \in \mathbb{R}^{\geq 0}$, a capacity $c_v \in \mathbb{R}^{\geq 0}$, and its available multiplicities $m_v \in \mathbb{Z}^{\geq 0}$.

By a demand assignment we mean a function $h: E \times V \rightarrow \mathbb{Z}^{\geq 0}$, where $h_{e,v}$ denotes the amount of demand that is assigned from edge e to vertex v . For any $v \in V$, we use $D_h(v)$ to denote the total amount of demand vertex v has received in h , i.e., $D_h(v) = \sum_{e \in E[v]} h_{e,v}$.

The corresponding multiplicity function, denoted $x^{(h)}$, is defined to be $x_v^{(h)} = \lceil D_h(v)/c_v \rceil$. A demand assignment h is *feasible* if $\sum_{v \in e} h_{e,v} \geq d_e$ for all $e \in E$ and $x_v^{(h)} \leq m_v$ for all $v \in V$. In other words, the demand of each edge is fully-assigned to (fully-served by) its incident vertices and the multiplicity of each vertex does not exceed its available multiplicities. The *weight (cost)* of h , denoted $w(h)$, is defined to be $\sum_{v \in V} w_v \cdot x_v^{(h)}$.

Given an instance $\Pi = (V, E, d_e, w_v, c_v, m_v)$ as described above, the problem of VC-HC is to compute a feasible demand assignment h such that $w(h)$ is minimized. Without loss of generality, we assume that the input graph G admits a feasible demand assignment.²

Augmented Cover.

Let $\Pi = (V, E, d_e, w_v, c_v, m_v)$ be an instance for VC-HC. For any integral $\beta \geq 1$, we say that a demand assignment h forms an augmented (β, γ) -cover if

1. $\sum_{v \in e} h_{e,v} \geq d_e$ for all $e \in E$.
2. $x_v^{(h)} \leq \beta \cdot m_v$ for all $v \in V$.
3. $w(h) \leq \gamma \cdot \min_{h' \in \mathcal{F}} w(h')$, where \mathcal{F} is the set of feasible demand assignments for Π .

2.2 LP Relaxation and the Dual LP

Let $\Pi = (V, E, d_e, w_v, c_v, m_v)$ be the input instance of VC-HC. The natural LP relaxation of VC-HC for the instance Π is given below in LP(1). The first three inequalities model the feasibility constraints of a demand assignment and its corresponding multiplicity function. The fourth inequality states that the multiplicity of a vertex cannot be zero if any demand gets assigned to it. This seemingly unnecessary constraint is required in giving a bounded integrality gap for this LP relaxation.

$$\begin{array}{ll}
 \text{Minimize} & \sum_{v \in V} w_v \cdot x_v & (1) \\
 \\
 \sum_{v \in e} h_{e,v} & \geq d_e, & \forall e \in E \\
 c_v \cdot x_v - \sum_{e \in E[v]} h_{e,v} & \geq 0, & \forall v \in V \\
 x_v & \leq m_v, & \forall v \in V \\
 d_e \cdot x_v - h_{e,v} & \geq 0, & \forall e \in E, v \in e \\
 x_v, h_{e,v} & \geq 0, & \forall e \in E, v \in e
 \end{array}$$

The dual LP for the instance Π is given below in LP(2). A solution $\Psi = (y_e, z_v, g_{e,v}, \eta_v)$ to this LP can be interpreted as an extended packing LP as follows: We want to raise the values of y_e for all $e \in E$. However, the value of each y_e is constrained by z_v and $g_{e,v}$ that are further constrained by w_v for each $v \in e$. The variable η_v provides an additional degree

² By selecting all of the available multiplicities, the feasibility of G can be checked via a max-flow computation.

of freedom in this packing program in that it allows higher values to be packed into y_e in the cost of a reduction in the objective value. Note that, this exchange does not always yield a better lower-bound for the optimal solution. In this paper we present an extended primal-dual scheme to handle this flexibility.

$$\begin{array}{ll}
 \text{Maximize} & \sum_{e \in E} d_e \cdot y_e - \sum_{v \in V} m_v \cdot \eta_v & (2) \\
 & c_v \cdot z_v + \sum_{e \in E[v]} d_e \cdot g_{e,v} - \eta_v \leq w_v, & \forall v \in V \\
 & y_e \leq z_v + g_{e,v}, & \forall v \in V, e \in E[v] \\
 & y_e, z_v, g_{e,v}, \eta_v \geq 0, & \forall v \in V, e \in E[v]
 \end{array}$$

For the rest of this paper, we will use $\text{OPT}(\Pi)$ to denote the cost of optimal solution for the instance Π . Since the optimal value of the above LPs gives a lower-bound on $\text{OPT}(\Pi)$ which we will be working with, we also use $\text{OPT}(\Pi)$ to denote their optimal value in the context.

3 A Primal-Dual Schema for VC-HC

In this section we present our extended primal-dual algorithm for VC-HC. The algorithm we present extends the framework developed for the soft capacity model [12, 7]. In the prior framework, the demand is assigned immediately when a vertex from its vicinity gets saturated. In our algorithm, we keep some of decisions pending until we have sufficient capacity for the demands. In contrast to the primal-dual scheme used in [6], which always stores dual values in $g_{e,v}$, we store the dual values in both $g_{e,v}$ and z_v , depending on the amount of unassigned demand v possesses in its vicinity. This ensures that, the cost of each multiplicity is charged only to the demands it serves.

To obtain a solid bound for this approach, however, we need to guarantee that the vertices whose multiplicity limits are attained receive sufficient amount of demands to charge to. This motivates our flow-based procedure Self-Containment for dealing with the pending decisions. During this procedure, a natural demand assignment is also formed.

3.1 The Algorithm

In this section we present our extended primal-dual algorithm DUAL-VCHC. This algorithm takes as input an instance $\Pi = (V, E, d, w, c, m)$ of VC-HC and outputs a feasible primal demand assignment h together with a feasible dual solution $\Psi = (y_v, z_v, g_{e,v}, \eta_v)$ for Π .

The algorithm starts with an initial zero dual solution and eventually reaches a locally optimal solution. During the process, the values of the dual variables in Ψ are raised gradually and some inequalities will meet with equality. We say that a vertex v is *saturated* if the inequality $c_v \cdot z_v + \sum_{e \in E[v]} d_e \cdot g_{e,v} - \eta_v \leq w_v$ is met with equality.

Let $E^\phi := \{e : e \in E, d_e > 0\}$ be the set of edges with non-zero demand and $V^\phi := \{v : v \in V, m_v \cdot c_v > 0\}$ be the set of vertices with non-zero capacity. For each $v \in V$, we use $d^\phi(v) = \sum_{e \in E[v] \cap E^\phi} d_e$ to denote the total amount of demand in $E[v] \cap E^\phi$. For intuition, E^ϕ contains the set of edges whose demands are not yet processed nor assigned, and V^ϕ corresponds to the set of vertices that have not yet saturated.

In addition, we maintain a set S , initialized to be empty, to denote the set of vertices that have saturated and that have at least one incident edge in E^ϕ . Intuitively, S corresponds to vertices with pending assignments.

The algorithm works as follows. Initially all dual variables in Ψ and the demand assignment h are set to be zero. We raise the value of the dual variable y_e for each $e \in E^\phi$ simultaneously at the same rate. To maintain the dual feasibility, as we increase y_e , either z_v or $g_{e,v}$ has to be raised for each $v \in e$. If $d^\phi(v) \leq c_v$, then we raise $g_{e,v}$. Otherwise, we raise z_v . In addition, for all $v \in e \cap S$, we raise η_v to the extent that keeps v saturated.

When a vertex $u \in V^\phi$ becomes saturated, it is removed from V^ϕ . Then we invoke a recursive procedure $\text{Self-Containment}(S \cup \{u\}, u)$, which we describe in the next paragraph, to compute a pair (S', h') , where

- S' is a maximal subset of $S \cup \{u\}$ whose capacity, if chosen, can fully-serve the demands in $E[S'] \cap E^\phi$, and
- h' is the corresponding demand assignment function (from $E[S'] \cap E^\phi$ to S').

If $S' = \emptyset$, then we leave the assignment decision pending and add u to S . Otherwise, S' is removed from S and $E[S']$ is removed from E^ϕ . In addition, we add the assignment h' to final assignment h to be output. This process repeats until $E^\phi = \emptyset$. Then the algorithm outputs h and Ψ and terminates.

We also note that, the particular vertex to saturate in each iteration is the one with the smallest value of $w^\phi(v) / \min\{c_v, d^\phi(v)\}$, where $w^\phi(v) := w_v - \left(c_v \cdot z_v + \sum_{e \in E[v]} d_e \cdot g_{e,v} - \eta_v\right)$ denotes the current slack of the inequality associated with $v \in V^\phi$.

The Procedure $\text{Self-Containment}(A, u)$

In the following we describe the recursive procedure $\text{Self-Containment}(A, u)$. It takes as input a vertex subset $A \subseteq V$ and a vertex $u \in V$ and outputs a pair (S', \tilde{h}') , where S' is a maximal subset of A whose capacity is sufficient to serve the unassigned demands in its vicinity, and h' is the corresponding demand assignment.

First we define a directed flow-graph $\mathcal{G}(A)$ with a source s^+ and a sink s^- for the vertex set A as follows. Excluding the source s^+ and the sink s^- , $\mathcal{G}(A)$ is a bipartite graph induced by $E[A] \cap E^\phi$ and A . For each $e \in E[A] \cap E^\phi$, we have a vertex \tilde{e} and an edge (s^+, \tilde{e}) in \mathcal{G} . Similarly, for each $v \in A$ we have a vertex \tilde{v} and an edge (\tilde{v}, s^-) . For each $v \in A$ and each $e \in E[v] \cap E^\phi$, we have an edge (\tilde{e}, \tilde{v}) in \mathcal{G} .

The capacity of each edge is defined as follows. For each $e \in E[A] \cap E^\phi$, the capacity of (s^+, \tilde{e}) is set to be d_e . For each $v \in A$, the capacity of (\tilde{v}, s^-) is set to be $m_v \cdot c_v$. The capacities of the remaining edges are unlimited.

The procedure Self-Containment works as follows. If $u \in A$, then it computes the max-flow \tilde{h} for $\mathcal{G}(A)$ with the additional constraint that $\tilde{h}(\tilde{u}, s^-)$ is minimized among all max-flows for $\mathcal{G}(A)$.³ If $u \notin A$, then it simply computes any max-flow \tilde{h} for $\mathcal{G}(A)$. Let

$$S' = \{v : v \in A \text{ such that } \tilde{h}(s^+, \tilde{e}) = d_e \text{ for all } e \in E[v] \cap E^\phi\}$$

be the subset of A that is able to serve the demand in $E[S'] \cap E^\phi$. If $S' = A$ or $S' = \emptyset$, then it returns (S', \tilde{h}') , where \tilde{h}' is the demand assignment induced by \tilde{h} . Otherwise it returns $\text{Self-Containment}(S', u)$.

³ This criterion can be achieved by imposing an additional constraint when computing the augmenting paths.

3.2 Properties of Dual-VCHC

Below we derive basic properties of our algorithm. Since the algorithm keeps the constraints feasible when increasing the dual variables, we know that Ψ is feasible for the dual LP for Π . In the following, we first show that h is a feasible demand assignment for Π as well. Then we derive properties we will be using when establishing the bi-approximation factor next section.

Feasibility of the demand assignment h

We begin with procedure Self-Containment. Let (S', \tilde{h}') be the pair returned by procedure Self-Containment($S \cup \{u\}, u$). The following lemma shows that S' is indeed maximal.

► **Lemma 2.** *If there exists a $B \subseteq S \cup \{u\}$ such that B can fully-serve the demand in $E[B] \cap E^\phi$, then $B \subseteq S'$.*

Proof. Let S_1, S_2, \dots, S_k , where $S_1 = S \cup \{u\} \supset S_2 \supset \dots \supset S_k = S'$, denote the input of the procedure Self-Containment($S \cup \{u\}, u$) in each recursion.

Below we argue that $B \subseteq S_i$ implies that $B \subseteq S_{i+1}$ for all $1 \leq i < k$. Let \tilde{h}_B denote a maximum flow for the flow graph $\mathcal{G}(B)$. Since B can fully-serve the demand in $E[B] \cap E^\phi$, we know that $\tilde{h}_B(s^+, \tilde{e}) = d_e$ for all $e \in E[B] \cap E^\phi$.

Consider the flow function computed by Maxflow($\mathcal{G}(S_i), u$) and denote it by \tilde{h}_i . If $\tilde{h}_i(s^+, \tilde{e}) < d_e$ for some $e \in E[B] \cap E^\phi$, then we embed \tilde{h}_B into \tilde{h}_i , i.e., cancel the flow from $E[B] \cap E^\phi$ to B in \tilde{h}_i and replace it by \tilde{h}_B . We see that the resulting flow strictly increases and remains valid for $\mathcal{G}(S_i)$, which is a contradiction to the fact that \tilde{h}_i is a maximum flow for $\mathcal{G}(S_i)$. Therefore, we know that $\tilde{h}_i(s^+, \tilde{e}) = d_e$ for all $e \in E[B] \cap E^\phi$ and the vertices of B must be included in S_{i+1} . This shows that $B \subseteq S_i$ for all $1 \leq i \leq k$. ◀

The following lemma states the feasibility of this primal-dual process.

► **Lemma 3.** *E^ϕ becomes empty in polynomial time. Furthermore, the assignments computed by Self-Containment during the process form a feasible demand assignment.*

The cost incurred by h

Below we consider the cost incurred by the partial assignments computed by Self-Containment. Let V_S denote the set of vertices that have been included in the set S . For any vertex v that has saturated, we use (S'_v, h'_v) to denote the particular pair returned by Self-Containment such that $v \in S'_v$. Note that, this pair (S'_v, h'_v) is uniquely defined for each v that has saturated. Therefore, we know that $h_{e,v} = (h'_v)_{e,v}$ holds for any $e \in E[v]$.

In the rest of this section, we will simply use $h_{e,v}$ when it refers to $(h'_v)_{e,v}$ for simplicity of notations. Recall that $D_{h'_v}(v)$ denotes the amount of demand v receives in h'_v . We have the following proposition for the dual solution $\Psi = (y_e, z_v, g_{e,v}, \eta_v)$, which follows directly from the way the dual variables are raised.


► **Proposition 4.** *For any $v \in V$ such that $d^\phi(v) > c_v$ when saturated, the following holds:*

- $z_v = y_e$ for all $e \in E[v]$ with $h_{e,v} > 0$.
- $\eta_v > 0$ only when $v \in V_S$.

The following lemma gives the properties for vertices in V_S .

► **Lemma 5.** *For any $v \in V_S$, we have*

1. $D_{h'_v}(v) = m_v \cdot c_v$.
2. $w_v \cdot m_v = D_{h'_v}(v) \cdot y_e - m_v \cdot \eta_v$ for all $e \in E[v]$ such that $h_{e,v} > 0$.



■ **Figure 1** Alternating paths in the flow-graph $\mathcal{G}(S')$.

Proof. First we prove that $D_{h'_v}(v) < m_v \cdot c_v$. Without loss of generality, we assume that $m_v \geq 1$ and $D_{h'_v}(v) < m_v \cdot c_v$ for a contradiction.

Consider the iteration for which the vertex v was removed from S and let u be the vertex that becomes saturated in that iteration. By Lemma 2, we know that in the beginning of that iteration, $\exists B \subseteq S$ such that B can fully-serve $E[B] \cap E^\phi$. Therefore it follows that $u \in S'_v$, for otherwise S'_v would have been removed from S in the previous iteration.

Consider the flow-graph $\mathcal{G}(S'_v)$ and the max-flow \tilde{h}'_v to which h'_v corresponds. We know that $\tilde{h}'_v(\tilde{e}, \tilde{u}) = 0$ for all $e \in E[v] \cap E^\phi$, for otherwise we have an alternating path $\tilde{u} \rightarrow \tilde{e} \rightarrow \tilde{v}$ so that we can reroute the flow $\tilde{e} \rightarrow \tilde{u} \rightarrow s^-$ to $\tilde{e} \rightarrow \tilde{v} \rightarrow s^-$, which is a contradiction to the fact that the max-flow we compute is the one that minimizes the flow from \tilde{u} to s^- .

Let $S_0 := \{v\}$ and $E_0 := E[v] \cap E^\phi$. For $i \geq 1$, consider the sets S_i and E_i defined as

$$S_i := \bigcup_{e \in E_{i-1}} \{v' : v' \in e \cap S'_v\} \quad \text{and} \quad E_i := E[S_i] \cap E^\phi.$$

Note that, $u \notin S_i$ implies that $S_i \subsetneq S_{i+1}$, for otherwise S_i would be a subset of S that can fully-serve $E[S_i] \cap E^\phi$ since the beginning of the iteration, a contradiction to Lemma 2. Therefore $u \in S_j$ for some $j \geq 1$ since $|S_i| \leq |S'_v| < \infty$. Let j_0 be the smallest integer such that $u \in S_{j_0}$. By definition we have $S_0 \subsetneq S_1 \subsetneq \dots \subsetneq S_{j_0} \subseteq S'_v$. This corresponds to an alternating path to which we can reroute the flow from u to v , a contradiction. See also Fig. 1 for an illustration. Therefore we have $D_{h'_v}(v) = m_v \cdot c_v$.

For the second half of this lemma, since $v \in V_S$, we know that $d^\phi(v) > c_v$ before it gets saturated. Therefore, by Proposition 4, we know that $y_e = z_v$ holds for all $e \in E[v]$ such that $h_{e,v} > 0$. It follows that $w_v = c_v \cdot z_v - \eta_v = c_v \cdot y_e - \eta_v$ and $w_v \cdot m_v = D_{h'_v}(v) \cdot y_e - m_v \cdot \eta_v$ as claimed. ◀

The following auxiliary lemma, which is carried over from the previous primal-dual framework, shows that, for any vertex v with $d^\phi(v) \leq c_v$ when saturated, we can locate at most c_v units of demands from $E[v]$ such that their dual value pays for w_v . This statement holds intuitively since v is saturated.

► **Lemma 6.** For any $v \in V$ with $d^\phi(v) \leq c_v$ when saturated, we can compute a function $\ell_v: E[v] \rightarrow \mathbb{R}^{\geq 0}$ such that the following holds:

(a) $0 \leq h_{e,v} \leq \ell_v(e) \leq d_e$, for all $e \in E[v]$.

(b) $\sum_{e \in E[v]} \ell_v(e) \leq c_v$.

(c) $\sum_{e \in E[v]} \ell_v(e) \cdot y_e = w_v$.

Intuitively, Proposition 4 and Lemma 5 provide a solid upper-bound for vertices whose capacity is fairly used. However, we remark that, this approach does not yield a solid guarantee for vertices whose capacity is barely used, i.e., $D_{h'_v}(v) \ll c_v$. The reason is that the demand that is served (charged) by vertices that have been included in S , i.e., those discussed in Lemma 5, cannot be charged again since their dual values are inflated during the primal-dual process.

4 Augmented Cover

In this section we establish the following theorem:

► **Theorem 7.** For any integer $k \geq 2$, we can compute an augmented $\left(k, \left(1 + \frac{1}{k-1}\right)(f-1)\right)$ -cover for VC-HC in polynomial time.

Let $\Pi = (V, E, d, w, c, m)$ be the input instance. Let $m'_v := k \cdot m_v$ denote the augmented multiplicity function for each $v \in V$. We invoke algorithm DUAL-VCHC on the instance $\Pi' = (V, E, d, w, c, m')$. Let h be the demand assignment and $\Psi = (y, z, g, \eta)$ be the dual solution output by the algorithm for Π' .

The following observation is crucial in establishing the bi-approximation ratio: The dual solution Ψ , which was computed for instance Π' , is also feasible for input instance Π .

► **Lemma 8.** Ψ is feasible for LP(2) with respect to Π . In other words, we have

$$\sum_{e \in E} d_e \cdot y_e - \sum_{v \in V} m_v \cdot \eta_v \leq OPT(\Pi).$$

Proof. The statements follow directly since LP(2) has the same feasible region for Π and Π' . ◀

It is also worth mentioning that, the assignment h computed by DUAL-VCHC already gives an augmented $\left(k, \left(1 + \frac{1}{k-1}\right)f\right)$ -cover. To obtain our claimed ratio, however, we further modify some of the demand assignments in h to achieve better utilization on the residue capacity of the vertices. Below we describe this procedure and establish the bi-approximation ratio.

Let V_S denote the set of vertices that have been included in S . For each $v \in V$ such that $D_h(v) < c_v$, let ℓ_v denote the function given by Lemma 6 with respect to v . We use h^* to denote the resulting assignment to obtain, where h^* is initialized to be h . For each $e \in E$, we repeat the following operation until no such vertex pair can be found:

■ Find a vertex pair $u \in e \setminus V_S$ and $v \in e$ such that

$$\begin{cases} h_{e,u}^* > 0, \\ D_{h^*}(u) > c_u, \end{cases} \quad \text{and} \quad \begin{cases} D_h(v) < c_v, \\ h_{e,v}^* < \ell_v(e). \end{cases}$$

Then reassign $\min \{ h_{e,u}^*, \ell_v(e) - h_{e,v}^* \}$ units of demand of e from u to v .

In particular, we set $\begin{cases} h_{e,u}^* = h_{e,u}^* - R_{u,v}, \\ h_{e,v}^* = h_{e,v}^* + R_{u,v}, \end{cases}$ where $R_{u,v} := \min \{ h_{e,u}^*, \ell_v(e) - h_{e,v}^* \}$.

40:10 $O(f)$ Bi-approximation for Capacitated Covering with Hard Capacities

Intuitively, in assignment h^* if some demand is currently assigned to a vertex in $V \setminus V_S$ that requires multiple multiplicities, then we try to reassign it to vertices that have surplus residue capacity (according to the function ℓ_v) to balance the load. Note that, in this process we do not use additional multiplicities of the vertices, and the reassignments are performed only between vertices not belonging to V_S .

The following lemma shows that, the cost incurred by vertices in $V \setminus V_S$ can be distributed to the dual variables of the edges.

► **Lemma 9.** *We have*

$$\sum_{v \in V \setminus V_S} w_v \cdot x_v^{(h^*)} \leq (f-1) \cdot \sum_{v \in V_S} \sum_{e \in E[v]} h_{e,v}^* \cdot y_e + f \cdot \sum_{v \in V \setminus V_S} \sum_{e \in E[v]} h_{e,v}^* \cdot y_e.$$

The following lemma provides a lower bound for $\text{OPT}(\Pi)$ in terms of the net sum of the dual values over the edges.

► **Lemma 10.** *We have*

$$\sum_{e \in E} d_e \cdot y_e \leq \frac{k}{k-1} \cdot \text{OPT}(\Pi).$$

Proof. For each $v \in V_S$, by Lemma 5 we have $\sum_{e \in E[v]} h_{e,v}^* = m'_v \cdot c_v = k \cdot m_v \cdot c_v$. Furthermore, by the way how η_v is raised, we know that $\eta_v \leq c_v \cdot z_v = c_v \cdot y_e$ holds for all $e \in E[v]$ such that $h_{e,v}^* > 0$. Therefore, it follows that

$$m_v \cdot \eta_v \leq m_v \cdot c_v \cdot y_e \leq \frac{1}{k} \cdot \sum_{e \in E[v]} h_{e,v}^* \cdot y_e. \quad (3)$$

By Inequality (3) and Lemma 8, it follows that

$$\sum_{e \in E} \left(d_e - \frac{1}{k} \cdot \sum_{v \in e \cap V_S} h_{e,v}^* \right) \cdot y_e \leq \text{OPT}(\Pi). \quad (4)$$

Therefore,

$$\begin{aligned} \sum_{e \in E} d_e \cdot y_e &= \sum_{v \in V} \sum_{e \in E[v]} h_{e,v}^* \cdot y_e \leq \sum_{e \in E} \left(\frac{k}{k-1} \cdot d_e - \frac{1}{k-1} \cdot \sum_{v \in e \cap V_S} h_{e,v}^* \right) \cdot y_e \\ &= \frac{k}{k-1} \cdot \left(\sum_{e \in E} \left(d_e - \frac{1}{k} \cdot \sum_{v \in V_S \cap e} h_{e,v}^* \right) \cdot y_e \right) \\ &\leq \frac{k}{k-1} \cdot \text{OPT}(\Pi), \end{aligned}$$

where the last inequality follows from Inequality (4). ◀

In the following we establish the bi-criteria approximation factor and prove Theorem 7.

► **Lemma 11.** *We have*

$$w(h^*) \leq \left(1 + \frac{1}{k-1} \right) \cdot (f-1) \cdot \text{OPT}(\Pi)$$

for any integer $k \geq 2$.

Proof. By Lemma 5, we have $D_h(v) = m'_v \cdot c_v = k \cdot m_v \cdot c_v$ for any $v \in V_S$. Therefore,

$$w_v \cdot x_v^{(h^*)} = (c_v \cdot z_v - \eta_v) \cdot k \cdot m_v = \sum_{e \in E[v]} h_{e,v}^* \cdot y_e - k \cdot m_v \cdot \eta_v.$$

Applying Lemma 9, we obtain

$$\begin{aligned} w(h^*) &= \sum_{v \in V_S} w_v \cdot x_v^{(h^*)} + \sum_{v \in (V \setminus V_S)} w_v \cdot x_v^{(h^*)} \\ &\leq \left(\sum_{v \in V_S} \sum_{e \in E[v]} h_{e,v}^* \cdot y_e - k \cdot \sum_{v \in V} m_v \cdot \eta_v \right) \\ &\quad + \left((f-1) \cdot \sum_{v \in V_S} \sum_{e \in E[v]} h_{e,v}^* \cdot y_e + f \cdot \sum_{v \in (V \setminus V_S)} \sum_{e \in E[v]} h_{e,v}^* \cdot y_e \right) \\ &= f \cdot \sum_{v \in V} \sum_{e \in E[v]} h_{e,v}^* \cdot y_e - k \cdot \sum_{v \in V} m_v \cdot \eta_v \\ &= k \cdot \left(\sum_{v \in V} \sum_{e \in E[v]} h_{e,v}^* \cdot y_e - \sum_{v \in V} m_v \cdot \eta_v \right) + (f-k) \cdot \sum_{v \in V} \sum_{e \in E[v]} h_{e,v}^* \cdot y_e. \end{aligned}$$

The former item is upper-bounded by $k \cdot \text{OPT}(\Pi)$ by Lemma 8. Combing the above with Lemma 10, we obtain

$$w(h^*) \leq \left(k + (f-k) \cdot \frac{k}{k-1} \right) \cdot \text{OPT}(\Pi) = \left(1 + \frac{1}{k-1} \right) \cdot (f-1) \cdot \text{OPT}(\Pi)$$

as claimed. ◀

5 Conclusion

We conclude with some future directions. In this paper we presented bi-approximations for augmented multiplicity constraints. It is also interesting to consider VC-HC with relaxed demand constraints, i.e., partial covers. The reduction framework for partial VC-HC provided by Cheung et al. [3] and the tight approximation for VC-HC provided by Kao [10] jointly provided an almost tight $f + \epsilon$ -approximation when the vertices are unweighted.

When the vertices are weighted, it is known that $O(\frac{1}{\epsilon})f$ bi-approximations can be obtained via simple LP rounding. Comparing to the $O(\frac{1}{\epsilon})$ bi-approximation result we can obtain for classical vertex cover, there is still a gap, and this would be an interesting direction to explore.

References

- 1 Reuven Bar-Yehuda and Shimon Even. A linear-time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms*, 2(2):198–203, 1981. doi:10.1016/0196-6774(81)90020-1.
- 2 Reuven Bar-Yehuda, Guy Flysher, Julián Mestre, and Dror 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 *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 Rajiv Gandhi, Eran Halperin, Samir Khuller, Guy Kortsarz, and Aravind Srinivasan. An improved approximation algorithm for vertex cover with hard capacities. *J. Comput. Syst. Sci.*, 72:16–33, February 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 Dorit S. Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on Computing*, 11(3):555–556, 1982. doi:10.1137/0211045.
- 9 Mong-Jen Kao. *An Algorithmic Approach to Local and Global Resource Allocations*. PhD thesis, National Taiwan University, 2012.
- 10 Mong-Jen Kao. Iterative partial rounding for vertex cover with hard capacities. manuscript, 2016.
- 11 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.
- 12 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.
- 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 Barna Saha and Samir Khuller. Set cover revisited: Hypergraph cover with hard capacities. In *ICALP'12*, pages 762–773, 2012. doi:10.1007/978-3-642-31594-7_64.
- 15 L. A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4):385–393, 1982. doi:10.1007/BF02579435.