


Optimal Distributed Covering Algorithms

Ran Ben-Basat 

Harvard University, Cambridge, MA, United States
ran@seas.harvard.edu

Guy Even 

Tel Aviv University, Tel Aviv, Israel
guy@eng.tau.ac.il

Ken-ichi Kawarabayashi 

National Institute of Informatics, Tokyo, Japan
k_keniti@nii.ac.jp

Gregory Schwartzman 

National Institute of Informatics, Tokyo, Japan
greg@nii.ac.jp

Abstract

We present a time-optimal deterministic distributed algorithm for approximating a minimum weight vertex cover in hypergraphs of rank f . This problem is equivalent to the Minimum Weight Set Cover problem in which the frequency of every element is bounded by f . The approximation factor of our algorithm is $(f + \varepsilon)$. Let Δ denote the maximum degree in the hypergraph. Our algorithm runs in the CONGEST model and requires $O(\log \Delta / \log \log \Delta)$ rounds, for constants $\varepsilon \in (0, 1]$ and $f \in \mathbb{N}^+$. This is the first distributed algorithm for this problem whose running time does not depend on the vertex weights nor the number of vertices. Thus adding another member to the exclusive family of *provably optimal* distributed algorithms.

For constant values of f and ε , our algorithm improves over the $(f + \varepsilon)$ -approximation algorithm of [18] whose running time is $O(\log \Delta + \log W)$, where W is the ratio between the largest and smallest vertex weights in the graph. Our algorithm also achieves an f -approximation for the problem in $O(f \log n)$ rounds, improving over the classical result of [15] that achieves a running time of $O(f \log^2 n)$. Finally, for weighted vertex cover ($f = 2$) our algorithm achieves a *deterministic* running time of $O(\log n)$, matching the *randomized* previously best result of [16].

We also show that integer covering-programs can be reduced to the Minimum Weight Set Cover problem in the distributed setting. This allows us to achieve an $(f + \varepsilon)$ -approximate integral solution in $O\left((1 + f/\log n) \cdot \left(\frac{\log \Delta}{\log \log \Delta} + (f \cdot \log M)^{1.01} \cdot \log \varepsilon^{-1} \cdot (\log \Delta)^{0.01}\right)\right)$ rounds, where f bounds the number of variables in a constraint, Δ bounds the number of constraints a variable appears in, and $M = \max\{1, \lceil 1/a_{\min} \rceil\}$, where a_{\min} is the smallest normalized constraint coefficient. This improves over the results of [18] for the integral case, which combined with rounding achieves the same guarantees in $O(\varepsilon^{-4} \cdot f^4 \cdot \log f \cdot \log(M \cdot \Delta))$ rounds.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases Distributed Algorithms, Approximation Algorithms, Vertex Cover, Set Cover

Digital Object Identifier 10.4230/LIPIcs.DISC.2019.5

Related Version <http://arxiv.org/abs/1902.09377>

Funding *Ran Ben-Basat*: Supported by the Zuckerman Institute, the Technion Hiroshi Fujiwara Cyber Security Research center, and the Israel Cyber Directorate.

Ken-ichi Kawarabayashi: Supported by JSPS Kakenhi Grant Number JP18H05291.

Gregory Schwartzman: Supported by JSPS Kakenhi Grant Number JP19K20216 and JP18H05291.

Acknowledgements We thank the anonymous reviewers for their helpful remarks.



© Ran Ben-Basat, Guy Even, Ken-ichi Kawarabayashi, and Gregory Schwartzman;
licensed under Creative Commons License CC-BY

33rd International Symposium on Distributed Computing (DISC 2019).

Editor: Jukka Suomela; Article No. 5; pp. 5:1–5:15



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

In the Minimum Weight Hypergraph Vertex Cover (MWHVC) problem, we are given a hypergraph $G = (V, E)$ with vertex weights $w : V \rightarrow \{1, \dots, W\}$.¹ The goal is to find a minimum weight cover $U \subseteq V$ such that $\forall e \in E : e \cap U \neq \emptyset$. In this paper we develop a distributed approximation algorithm for MWHVC in the CONGEST model. The approximation ratio is $f + \varepsilon$, where f denotes the rank of the hypergraph (i.e., f is an upper on the size of every hyperedge). The MWHVC problem is a generalization of the Minimum Weight Vertex Cover (MWVC) problem (in which $f = 2$). The MWHVC problem is also equivalent to the Minimum Weight Set Cover Problem (the rank f of the hypergraph corresponds to the maximum frequency of an element). Both of these problems are among the classical NP-hard problems presented in [14].

We consider the following distributed setting for the MWHVC problem. The communication network is a bipartite graph $H(E \cup V, \{\{e, v\} \mid v \in e\})$. We refer to the network vertices as *nodes* and network edges as *links*. The nodes of the network are the hypergraph vertices on one side and hyperedges on the other side. There is a network link between vertex $v \in V$ and hyperedge $e \in E$ iff $v \in e$. The computation is performed in synchronous rounds, where messages are sent between neighbors in the communication network. As for message size, we consider the CONGEST model where message sizes are bounded to $O(\log |V|)$. This is more restrictive than the LOCAL model where message sizes are unbounded.

1.1 Related work

We survey previous results for MWHVC and MWVC. A comprehensive list of previous results appears in Tables 1 and Table 2.

Vertex Cover. The understanding of the round complexity for distributed MWVC has been established in two papers: a lower bound in [19] and a matching upper bound in [4]. Let Δ denote the maximum vertex degree in the graph G . The lower bound states that any distributed constant-factor approximation algorithm requires $\Omega(\log \Delta / \log \log \Delta)$ rounds to terminate. This lower bound holds for every constant approximation ratio, for unweighted graphs and even if the message lengths are not bounded (i.e., LOCAL model) [19]. The matching upper bound is a $(2 + \varepsilon)$ -approximation distributed algorithm in the CONGEST model, for every $\varepsilon = \Omega(\log \log \Delta / \log \Delta)$.² In [16] an $O(\log n)$ -round 2-approximation randomized algorithm for weighted graphs in the CONGEST model is given. We note that [16] was the first to achieve this running time with no dependence on W , the maximum weight of the nodes. Recently, Ben-Basat et al. [6] showed an $O(OPT^2 \log OPT)$ rounds algorithm for computing the minimal (unweighted) vertex cover and a $O(OPT)$ rounds for a $(2 + \varepsilon)$ -approximation. Here, OPT is the size of the smallest cover and thus these algorithms are adequate when a small solution exists.

Hypergraph Weighted Vertex Cover. For constant values of f , Astrand et al. [2] present an f -approximation algorithm for anonymous networks whose running time is $O(\Delta^2 + \Delta \cdot \log^* W)$. Khuller et al. [15] provide a solution that runs in $O(f \cdot \log \varepsilon^{-1} \cdot \log n)$ rounds in the CONGEST model for any $\varepsilon > 0$ and achieves an $(f + \varepsilon)$ -approximation. Setting $\varepsilon = 1/W$ (recall that

¹ Let $n \triangleq |V|$. We assume that $|E| = n^{O(1)}$ and $W = n^{O(1)}$.

² Recently, the range of ε for which the runtime is optimal was improved to $\Omega(\log^{-c} \Delta)$ for any $c = O(1)$ [5].

$W = \text{poly}(n)$) results in a f -approximation in $O(f \log^2 n)$ -rounds. For constant ε and f values, Kuhn et al. [17, 18] present an $(f + \varepsilon)$ -approximation algorithm that terminates in $O(\log \Delta + \log W)$ rounds.

For the Minimum Cardinality (i.e., unweighted) Vertex Cover in Hypergraphs Problem, the lower bound was recently matched by [9] with an $(f + \varepsilon)$ -approximation algorithm in the CONGEST model. The round complexity in [9] is $O\left(f/\varepsilon \cdot \frac{\log(f \cdot \Delta)}{\log \log(f \cdot \Delta)}\right)$, which is optimal for constant f and ε .³ The algorithm in [9] and its analysis is a deterministic version of the randomized maximal independent set algorithm of [10].

1.2 Our contributions

In this paper, we present a deterministic distributed $(f + \varepsilon)$ -approximation algorithm for minimum weight vertex cover in f -rank hypergraphs, which completes in

$$O\left(f \cdot \log(f/\varepsilon) + \frac{\log \Delta}{\log \log \Delta} + \min\{\log \Delta, f \cdot \log(f/\varepsilon) \cdot (\log \Delta)^{0.001}\}\right)$$

rounds in the CONGEST model. For any constants $\varepsilon \in (0, 1)$ and $f \in \mathbb{N}^+$ this implies a running time of $O(\log \Delta / \log \log \Delta)$, which is optimal according to [19]. This is the first distributed algorithm for this problem whose round complexity does not depend on the node weights nor the number of vertices.

Our algorithm is one of a handful of distributed approximation algorithms for *local* problems which are *provably optimal* [8, 7, 3, 4, 11, 9]. Among these are the classic Cole-Vishkin algorithm [8] for 3-coloring a ring, the more recent results of [3] and [4] for MWVC and Maximum Matching, and the result of [9] for Minimum Cardinality Hypergraph Vertex Cover.

Our algorithm also achieves a *deterministic* f -approximation for the problem in $O(f \log n)$ rounds. This improves over the best known deterministic result for hypergraphs $O(f \log^2 n)$ [15] and matches the best known *randomized* results for weighted vertex cover ($f = 2$) of $O(\log n)$ -rounds [16].

We also show that general covering Integer Linear Programs (ILPs) can be reduced to MWHVC in the distributed setting. That is, LP constraints can be translated into hyperedges such that a cover for the hyperedges satisfies all covering constraints. This allows us to achieve an $O\left((1 + f/\log n) \cdot \left(\frac{\log \Delta}{\log \log \Delta} + (f \cdot \log M)^{1.01} \cdot \log \varepsilon^{-1} \cdot (\log \Delta)^{0.01}\right)\right)$ rounds $(f + \varepsilon)$ -approximate *integral* solution, where f bounds the number of variables in a constraint, Δ bounds the number of constraints a variable appears in, and $M = \max\{1, \lceil 1/a_{\min} \rceil\}$, where a_{\min} is the smallest normalized constraint coefficient. This significantly improves over the results of [18] for the integral case, which combined with rounding achieves the same guarantees in $O(\varepsilon^{-4} \cdot f^4 \cdot \log f \cdot \log(M \cdot \Delta))$ rounds. Note that the results of [18] also include a $(1 + \varepsilon)$ -approximation for the fractional case, while our result only allows for an integral solution. We also note that plugging $\varepsilon = 1/(nWM)$ into our algorithm, achieves an f -approximation for ILPs in polylogarithmic time, a similar result cannot be achieved using [18].

³ The authors state their result for an $f(1 + \varepsilon)$ -approximation which removes the f factor from the runtime.

5:4 Optimal Distributed Covering Algorithms

■ **Table 1** Previous distributed algorithms for MWVC. In the table, $n = |V|$ and $\varepsilon \in (0, 1)$. Some of the algorithms hold only for the unweighted case and some are randomized. For randomized algorithms the running times hold in expectation or with high probability.

det.	weighted	approximation	time	algorithm
yes	no	3	$O(\Delta)$	[21]
yes	no	2	$O(\Delta^2)$	[1]
yes	yes	2	$O(1)$ for $\Delta \leq 3$	[1]
yes	yes	2	$O(\Delta + \log^* n)$	[20]
yes	yes	2	$O(\Delta + \log^* W)$	[2]
yes	yes	2	$O(\log^2 n)$	[15]
yes	yes	2	$O(\log n \log \Delta / \log^2 \log \Delta)$	[5]
no	yes	2	$O(\log n)$	[12, 16]
yes	yes	2	$O(\log n)$	This work
yes	yes	$2 + \varepsilon$	$O(\varepsilon^{-4} \log(W \cdot \Delta))$	[13, 18]
yes	yes	$2 + \varepsilon$	$O(\log \varepsilon^{-1} \log n)$	[15]
yes	yes	$2 + \varepsilon$	$O(\varepsilon^{-1} \log \Delta / \log \log \Delta)$	[4]
yes	yes	$2 + \varepsilon$	$O\left(\frac{\log \Delta}{\log \log \Delta} + \frac{\log \varepsilon^{-1} \log \Delta}{\log^2 \log \Delta}\right)$	[5]
yes	yes	$2 + \varepsilon$	$O\left(\frac{\log \Delta}{\log \log \Delta} + \log \varepsilon^{-1} \cdot (\log \Delta)^{0.001}\right)$	This work
yes	yes	$2 + \frac{\log \log \Delta}{c \cdot \log \Delta}$	$O(\log \Delta / \log \log \Delta)$	[4], $\forall c = O(1)$
yes	yes	$2 + (\log \Delta)^{-c}$	$O(\log \Delta / \log \log \Delta)$	[5], $\forall c = O(1)$
yes	yes	$2 + 2^{-c \cdot (\log \Delta)^{0.99}}$	$O(\log \Delta / \log \log \Delta)$	This work, $\forall c = O(1)$

■ **Table 2** Previous distributed algorithms for MWHVC. In the table, $n = |V|$ and $\varepsilon \in (0, 1)$. All algorithms are deterministic. Note that [9] holds only for unweighted hypergraphs.

weighted	approximation	time	algorithm
yes	f	$O(f^2 \Delta^2 + f \Delta \log^* W)$	[2]
yes	f	$O(f \log^2 n)$	[15]
yes	f	$O(f \log n)$	This work
no	$f + \varepsilon$	$O\left(\varepsilon^{-1} \cdot f \cdot \frac{\log(f \Delta)}{\log \log(f \Delta)}\right)$	[9]
yes	$f + \varepsilon$	$O(f \cdot \log(f/\varepsilon) \cdot \log n)$	[15]
yes	$f + \varepsilon$	$O(\varepsilon^{-4} \cdot f^4 \cdot \log f \cdot \log(W \cdot \Delta))$	[18]
yes	$f + \varepsilon$	$O\left(f \cdot \log(f/\varepsilon) \cdot (\log \Delta)^{0.001} + \frac{\log \Delta}{\log \log \Delta}\right)$	This work
no	$f + 1/c$	$O(\log \Delta / \log \log \Delta)$	[9], $\forall f, c = O(1)$
yes	$f + 2^{-c \cdot (\log \Delta)^{0.99}}$	$O(\log \Delta / \log \log \Delta)$	This work, $\forall f, c = O(1)$

1.3 Tools and techniques

The Primal-Dual schema. The Primal-Dual approach introduces, for every hyperedge $e \in E$, a dual variable denoted by $\delta(e)$. The dual edge packing constraints are $\forall v \in V, \sum_{v \in e} \delta(e) \leq w(v)$. If for some $\beta \in [0, 1)$ it holds that $\sum_{v \in e} \delta(e) \geq (1 - \beta) \cdot w(v)$, we say the node v is β -tight. Let $\beta = \varepsilon / (f + \varepsilon)$. For every feasible dual solution, the weight of the set of β -tight vertices is at most $(f + \varepsilon)$ times the weight of an optimal (fractional) solution. The algorithm terminates when the set of β -tight edges constitutes a vertex cover.

The challenge. When designing a Primal-Dual distributed algorithm, the main challenge is in controlling the rate at which we increase the dual variables. On the one hand, we must grow them rapidly to reduce the number of communication rounds. On the other hand, we may not violate the edge packing constraints. This is tricky in the distributed environments as we have to coordinate between nodes. For example, the result of [4] does not generalize to hypergraphs, as hyperedges require the coordination of more than two nodes in order to increment edge variables.

Our algorithm. The algorithm proceeds in iterations, each of which requires a constant number of communication rounds. We initialize the dual variables in a “safe” way so that feasibility is guaranteed. We refer to the additive increase of the dual variable $\delta(e)$ as $\text{bid}(e)$. Similarly to [5] (where bids are called deals), we use *levels* to measure the progress made by a vertex. Whenever the level of a vertex increases, it sends a message about it to all incident edges, which multiply (decrease) their bids by 0.5. Intuitively, the level of a vertex equals the logarithm of its uncovered portion. Formally, we define the level of a vertex v as $\ell(v) \triangleq \left\lceil \log \frac{w(v)}{w(v) - \sum_{e \ni v} \delta(e)} \right\rceil$. That is, the initial level of v is 0 and it is increased as the dual variables of the edges incident to v grow. The level of a vertex never reaches $z \triangleq \lceil \log \beta^{-1} \rceil$ as this implies that it is β -tight and entered the cover. Loosely speaking, the algorithm increases the increments $\text{bid}(e)$ exponentially (multiplication by α) provided that no vertex $v \in e$ is $(0.5^{\ell(v)}/\alpha)$ -tight with respect to the bids of the previous iteration. Here, $\alpha \geq 2$ is a positive parameter that we determine later. The analysis builds on two observations: (1) The number of times that the increment $\text{bid}(e)$ is multiplied by α is bounded by $\log_\alpha \Delta$. (2) The number of iterations in which $\text{bid}(e)$ is not multiplied by α is bounded by $O(f \cdot z \cdot \alpha)$. Loosely speaking, each such iteration means that for at least one vertex $v \in e$ the sum of bids is at least an $1/(2\alpha)$ -fraction of its slack. Therefore, after at most $O(\alpha)$ such iterations that vertex will level up. Since there are z levels per vertex and f vertices in e , we have that the number of such iterations is at most $O(f \cdot z \cdot \alpha)$. Hence the total number of iterations is bounded by $O(\log_\alpha \Delta + f \cdot z \cdot \alpha)$.

Integer linear programs (ILPs). We show distributed reductions that allow us to compute an $(f + \varepsilon)$ -approximation for general covering integer linear programs (ILPs). To that end, we first show that any Zero-One covering program (where all variables are binary) can be translated into a set cover instance in which the vertex degree is bounded by 2^f times the bound on the number of constraints each ILP variable appears in. We then generalize to arbitrary covering ILPs by replacing each variable with multiple vertices in the hypergraph, such that the value associated with the ILP variable will be the weighted sum of the vertices in the cover.

2 Problem Formulation

Let $G = (V, E)$ denote a hypergraph. Vertices in V are equipped with positive weights $w(v)$. For a subset $U \subseteq V$, let $w(U) \triangleq \sum_{v \in U} w(v)$. Let $E(U)$ denote the set of hyperedges that are incident to some vertex in U (i.e., $E(U) \triangleq \{e \in E \mid e \cap U \neq \emptyset\}$).

The *Minimum Weight Hypergraph Vertex Cover* Problem (MWHVC) is defined as follows.

- Input:** Hypergraph $G = (V, E)$ with vertex weights $w(v)$.
Output: A subset $C \subseteq V$ such that $E(C) = E$.
Objective: Minimize $w(C)$.

The MWHVC Problem is equivalent to the Weighted Set Cover Problem. Consider a set system (X, \mathcal{U}) , where X denotes a set of elements and $\mathcal{U} = \{U_1, \dots, U_m\}$ denotes a collection of subsets of X . The reduction from the set system (X, \mathcal{U}) to a hypergraph $G = (V, E)$ proceeds as follows. The set of vertices is $V \triangleq \{u_1, \dots, u_m\}$ (one vertex u_i per subset U_i). The set of edges is $E \triangleq \{e_x\}_{x \in X}$ (one hyperedge e_x per element x), where $e_x \triangleq \{u_i : x \in U_i\}$. The weight of vertex u_i equals the weight of the subset U_i . We now detail the setting for computing a distributed $(f + \varepsilon)$ -approximation of the problem.

5:6 Optimal Distributed Covering Algorithms

Input. The input is a hypergraph $G = (V, E)$ with non-negative vertex weights $w : V \rightarrow \mathbb{N}^+$ and an approximation ratio parameter $\varepsilon \in (0, 1]$. We denote the number of vertices by n , the rank of G by f (i.e., each hyperedge contains at most f vertices), and the maximum degree of G by Δ (i.e., each vertex belongs to at most Δ edges).

Assumptions. We assume that

- (i) Vertex weights are polynomial in $n \triangleq |V|$ so that sending a vertex weight requires $O(\log n)$ bits.
- (ii) Vertex degrees are polynomial in n (i.e., $|E(v)| = n^{O(1)}$) so that sending a vertex degree requires $O(\log n)$ bits. Since $|E(v)| \leq n^f$, this assumption trivially holds for constant f .
- (iii) The maximum degree is at least 3 so that $\log \log \Delta > 0$.

Output. The nodes compute a vertex cover $C \subseteq V$. Namely, for every hyperedge $e \in E$, the intersection $e \cap C$ is not empty. The set C is maintained locally in the sense that every vertex v knows whether it belongs to C or not.

Communication Network. The communication network $N(E \cup V, \{\{e, v\} \mid v \in e\})$ is a bipartite graph. There are two types of nodes in the network: *servers* and *clients*. The set of servers is V (the vertex set of G) and the set of clients is E (the hyperedges in G). There is a link (v, e) from server $v \in V$ to a client $e \in E$ if $v \in e$. We note that the degree of the clients is bounded by f and the degree of the servers is bounded by Δ .

Notation.

- We say that an edge e is *covered* by C if $e \cap C \neq \emptyset$.
- Let $E(v) \triangleq \{e \in E \mid v \in e\}$ denote the set of hyperedges that contain v .
- For every vertex v , the algorithm maintains a subset $E'(v) \subseteq E(v)$ that consists of the uncovered hyperedges in $E(v)$ (i.e., $E'(v) = \{e \in E(v) \mid e \cap C = \emptyset\}$).

3 Distributed Approximation Algorithm for MWHVC

3.1 Parameters and Variables

- The algorithm computes an $(f + \varepsilon)$ -approximation where $\varepsilon \in (0, 1]$. The parameter β is defined by $\beta \triangleq \varepsilon / (f + \varepsilon)$, where f is the rank of the hypergraph.
- Each vertex v is assigned a level $\ell(v)$ which is a nonnegative integer.
- We denote the dual variables at the end of iteration i by $\delta_i(e)$ (see Appendix A for a description of the dual edge packing linear program). The amount by which $\delta_i(e)$ is increased in iteration i is denoted by $\text{bid}_i(e)$. Namely, $\delta_i(e) = \sum_{j \leq i} \text{bid}_j(e)$.
- The parameter $\alpha \geq 2$ determines the factor by which bids are multiplied. We determine its value in the analysis in the following section.

3.2 Algorithm MWHVC

1. Initialization. Set $C \leftarrow \emptyset$. For every vertex v , set level $\ell(v) \leftarrow 0$ and uncovered edges $E'(v) \leftarrow E(v)$.
2. Iteration $i = 0$. Every edge e collects the weight $w(v)$ and degree $|E(v)|$ from every vertex $v \in e$, and sets: $\text{bid}(e) = 0.5 \cdot \min_{v \in e} \{w(v) / |E(v)|\}$. The value $\text{bid}(e)$ is sent to every $v \in e$. The dual variable is set to $\delta(e) \leftarrow \text{bid}(e)$.

3. For $i = 1$ to ∞ do:
 - a. Check β -tightness. For every $v \notin C$, if $\sum_{e \in E(v)} \delta(e) \geq (1 - \beta)w(v)$, then v joins the cover C , sends a message to every $e \in E'(v)$ that e is covered, and vertex v terminates.
 - b. For every uncovered edge e , if e receives a message that it is covered, then it tells all its vertices that e is covered and terminates.
 - c. For every vertex $v \notin C$, if it receives a message from e that e is covered, then $E'(v) \leftarrow E'(v) \setminus \{e\}$. If $E'(v) = \emptyset$, then v terminates (without joining the cover).
 - d. Increment levels and scale bids.
 For every active (that has not terminated) vertex⁴:
 While $\sum_{e \in E(v)} \delta(e) > w(v)(1 - 0.5^{\ell(v)+1})$ do
 - i. $\ell(v) \leftarrow \ell(v) + 1$
 - ii. For every $e \in E'(v)$: $\text{bid}(e) \leftarrow 0.5 \cdot \text{bid}(e)$
 - e. For every active vertex, if $\sum_{e \in E'(v)} \text{bid}(e) \leq \frac{1}{\alpha} \cdot 0.5^{\ell(v)+1} \cdot w(v)$, then send the message “raise” to every $e \in E'(v)$; otherwise, send the message “stuck” to every $e \in E'(v)$.
 - f. For every uncovered edge e , if *all* incoming messages are “raise” $\text{bid}(e) \leftarrow \alpha \cdot \text{bid}(e)$. Send $\text{bid}(e)$ to every $v \in e$, who updates $\delta(e) \leftarrow \delta(e) + \text{bid}(e)$.
- **Termination.** Every vertex v terminates when either $v \in C$ or every edge $e \in E(v)$ is covered (i.e., $E'(v) = \emptyset$). Every edge e terminates when it is covered (i.e., $e \cap C \neq \emptyset$). We say that the algorithm has terminated if all the vertices and edges have terminated.
- **Execution in congest.** See Section B in the Appendix for a discussion of how Algorithm MWHVC is executed in the CONGEST model.

4 Algorithm Analysis

In this section, we analyze the approximation ratio and the running time of the algorithm. Throughout the analysis, we attach an index i to the variables $\text{bid}_i(e)$, $\delta_i(e)$ and $\ell_i(v)$. The indexed variable refers to its value at the end of the i 'th iteration.

4.1 Feasibility and Approximation Ratio

The following invariants are satisfied throughout the execution of the algorithm. In the following claim we bound the sum of the bids of edges incident to a vertex.

▷ **Claim 1.** If $v \notin C$ at the end of iteration i , then $\sum_{e \in E'(v)} \text{bid}_i(e) \leq 0.5^{\ell_i(v)+1} \cdot w(v)$.

Proof. The proof is by induction on i . For $i = 0$, the claim holds because $\ell_0(v) = 0$ and $\text{bid}_0(e) \leq 0.5 \cdot w(v)/|E(v)|$. The induction step, for $i \geq 1$, considers two cases. (A) If v sends a “raise” message in iteration i , then Step 3e implies that $\sum_{e \in E'(v)} \text{bid}_i(e) \leq 0.5^{\ell(v)+1} \cdot w(v)$, as required. (B) Suppose v sends a “stuck” message in iteration i . By Step 3d, $\text{bid}_i(e) \leq 0.5^{\ell_i(v) - \ell_{i-1}(v)} \cdot \text{bid}_{i-1}(e)$ for every $e \in E'(v)$. The induction hypothesis states that $\sum_{e \in E'(v)} \text{bid}_{i-1}(e) \leq 0.5^{\ell_{i-1}(v)+1} \cdot w(v)$. The claim follows by combining these inequalities. ◁

If an edge e is covered in iteration j , then e terminates and $\delta_i(e)$ is not set for $i \geq j$. In this case, we define $\delta_i(e) = \delta_{j-1}(e)$, namely, the last value assigned to a dual variable.

⁴ For simplicity of the description, we assume in step 3(d)ii that every $v \in e$ decides whether $\text{bid}(e)$ is halved. In a distributed setting, $\text{bid}(e)$ is halved if there exists a vertex $v \in e$ that requests such a halving. The distributed implementation of this step is further discussed in Appendix B.

5:8 Optimal Distributed Covering Algorithms

▷ **Claim 2.** For every $i \geq 1$ and every a vertex v that has not terminated the following inequality holds:

$$w(v)(1 - 0.5^{\ell_i(v)}) \leq \sum_{e \in E(v)} \delta_{i-1}(e) \leq (1 - 0.5^{\ell_i(v)+1}) \cdot w(v). \quad (1)$$

In addition the dual variables $\delta_i(e)$ constitute a feasible edge packing. Namely,

$$\begin{aligned} \sum_{e \in E(v)} \delta_i(e) &\leq w(v) && \text{for every vertex } v \in V, \\ \delta_i(e) &\geq 0 && \text{for every edge } e \in E. \end{aligned}$$

Proof. We prove the claim by induction on the iteration number i . To simplify the proof, we reformulate the statement of the feasibility of the dual variables to $i - 1$. We first prove the induction basis for $i = 1$.

Proof of Eq. 1 for $i = 1$. Fix a vertex v . At the end of iteration 0, $\ell_0(v) = 0$ and $0 < \text{bid}_0(e) \leq w(v)/(2|E(v)|)$, for every $e \in E(v)$. Hence $0 < \sum_{e \in E(v)} \text{bid}_0(e) \leq w(v)/2$. Because $\delta_0(e) = \text{bid}_0(e)$, the condition in Step 3d does not hold, and $\ell_1(v) = \ell_0(v) = 0$. We conclude that Eq. 1 holds for $i = 1$.

Proof of feasibility of $\delta_0(e)$ for $i = 1$. Non-negativity follows from the fact that $\delta_0(e) = \text{bid}_0(e) > 0$. The packing constraint for vertex v is satisfied because $\sum_{e \in E(v)} \text{bid}_0(e) \leq w(v)/2$. This completes the proof of the induction basis.

We now prove the induction step assuming that Eq. 1 holds for δ_{i-1} .

Proof of Eq. 1 for $i > 1$. Since v is not in the cover it is also not β -tight. Step 3d in iteration i increases $\ell(v)$ until Eq. 1 holds for i .

Proof of feasibility of $\delta_{i-1}(e)$ for $i > 1$. Consider a vertex v . If v joins C in iteration $i - 1$, then $\delta_{i-1} = \delta_{i-2}$, and the packing constraint of v holds by the induction hypothesis. Otherwise, then by $\delta_{i-1}(e) = \delta_{i-2}(e) + \text{bid}_{i-1}(e)$, Claim 1, and the induction hypothesis for Eq. 1, we have

$$\begin{aligned} \sum_{e \in E(v)} \delta_{i-1}(e) &= \sum_{e \in E(v)} (\delta_{i-2}(e) + \text{bid}_{i-1}(e)) \\ &\leq \left(1 - 0.5^{\ell_{i-1}(v)+1} + 0.5^{\ell_{i-1}(v)+1}\right) \cdot w(v) = w(v). \quad \triangleleft \end{aligned}$$

Let opt denote the cost of an optimal (fractional) weighted vertex cover of G .

► **Corollary 3.** Upon termination, the approximation ratio of Algorithm MWHVC is $f + \varepsilon$.

Proof. Throughout the algorithm, the set C consists of β -tight vertices. By Claim 20, $w(C) \leq (f + \varepsilon) \cdot \text{opt}$. Upon termination, C constitutes a vertex cover (as an edge only terminate once covered), and the corollary follows. ◀

4.2 Communication Rounds Analysis

In this section, we prove that the number of communication rounds of Algorithm MWHVC is bounded by (where $\gamma > 0$ is a constant, e.g., $\gamma = 0.001$)

$$O\left(f \cdot \log(f/\varepsilon) + \frac{\log \Delta}{\gamma \cdot \log \log \Delta} + \min\{\log \Delta, f \cdot \log(f/\varepsilon) \cdot (\log \Delta)^\gamma\}\right).$$

It suffices to bound the number of iterations because each iteration consists of a constant number of communication rounds.

Let $z \triangleq \lceil \log_2 \frac{1}{\beta} \rceil$. Note that $z = O(\log(f/\varepsilon))$.

▷ Claim 4. The level of every vertex is always less than z .

Proof. Assume that $\ell_i(v) \geq z$. By Eq. 1, $\sum_{e \in E(v)} \delta_{i-1}(e) \geq w(v) \cdot (1 - 2^{-z}) \geq (1 - \beta) \cdot w(v)$. This implies that v is β -tight and joins the cover in Line 3a before $\ell_i(v)$ reaches z . ◀

4.2.1 Raise or Stuck Iterations

► **Definition 5** (*e*-raise and *v*-stuck iterations). *An iteration $i \geq 1$ is an e-raise iteration if in Line 3f we multiplied $\text{bid}(e)$ by α . An iteration $i \geq 1$ is a v-stuck iteration if v sent the message “stuck” in iteration i .*

Note that if iteration i is a *v*-stuck iteration and $v \in e$, then $\text{bid}_i(e) \leq \text{bid}_{i-1}(e)$ and i is not an *e*-raise iteration.

We bound the number of *e*-raise iterations as follows.

► **Lemma 6.** *The number of e-raise iterations is bounded by $\log_\alpha(\Delta \cdot 2^{f \cdot z})$.*

Proof. Let v^* denote a vertex with minimum normalized weight in e (that is, $v^* \in \arg\min_{v \in e} \{w(v)/|E(v)|\}$). The first bid satisfies $\text{bid}_0(e) = 0.5 \cdot w(v^*)/|E(v^*)| \geq 0.5 \cdot w(v^*)/\Delta$. By Claim 1, $\text{bid}_i(e) \leq 0.5 \cdot w(v^*)$. The bid is multiplied by α in each *e*-raise iteration and is halved at most $f \cdot z$ times. The bound on the number of halvings holds because the number of vertices in the edge is bounded by f , and each vertex halves the bid each time its level is incremented. The lemma follows. ◀

We bound the number of *v*-stuck iterations as follows.

► **Lemma 7.** *For every vertex v and level $\ell(v)$, the number of v-stuck iterations is bounded by α .*

Proof. Notice that when v reached the level $\ell(v)$, we had $\sum_{e \in E(v)} \delta(e) \geq w(v)(1 - 0.5^{\ell(v)})$. The number of *v*-stuck iterations is then bounded by the number of times it can send a “stuck” message without reaching $\sum_{e \in E(v)} \delta(e) > w(v)(1 - 0.5^{\ell(v)+1})$. Indeed, once this inequality holds, the level of v is incremented. Every stuck iteration implies, by Line 3e, that $\sum_{e \in E'(v)} \text{bid}(e) > \frac{1}{\alpha} \cdot 0.5^{\ell(v)+1} \cdot w(v)$. Therefore, we can bound the number of iteration by $\frac{w(v)(1 - 0.5^{\ell(v)+1}) - w(v)(1 - 0.5^{\ell(v)})}{\frac{1}{\alpha} \cdot 0.5^{\ell(v)+1} \cdot w(v)} = \alpha$. ◀

4.2.2 Bound on the Number of Rounds

► **Theorem 8.** *For every $\alpha \geq 2$, the number of iterations of Algorithm MWHVC is*

$$O(\log_\alpha(\Delta \cdot 2^{f \cdot z}) + f \cdot z \cdot \alpha) = O(\log_\alpha \Delta + f \cdot \log(f/\varepsilon) \cdot \alpha).$$

Proof. Fix an edge e . We bound the number of iterations until e is covered. Every iteration is either an *e*-raise iteration or a *v*-stuck iteration for some $v \in e$. Since e contains at most f vertices, we conclude that the number of iterations is bounded by the number of *e*-raise iterations plus the sum over $v \in e$ of the number of *v*-stuck iterations. The theorem follows from Lemmas 6 and 7. ◀

In Theorem 9 we assume that all the vertices know the maximum degree Δ and that $\Delta \geq 3$. The assumption that the maximal degree Δ is known to all vertices is not required. Instead, each hyperedge e can compute a local maximum degree $\Delta(e)$, where $\Delta(e) \triangleq \max_{u \in e} |E(u)|$. The local maximum degree $\Delta(e)$ can be used instead of Δ to define local value of the multiplier $\alpha = \alpha(e)$. Let $T(f, \Delta, \varepsilon)$ denote the round complexity of Algorithm MWHVC. By setting α appropriately, we bound the running time as follows.

5:10 Optimal Distributed Covering Algorithms

► **Theorem 9.**⁵ Let $\gamma > 0$ denote a constant and set

$$\alpha = \begin{cases} \max\left(2, \frac{\log \Delta}{f \cdot \log(f/\varepsilon) \cdot \log \log \Delta}\right) & \text{if } \frac{\log \Delta}{f \cdot \log(f/\varepsilon) \cdot \log \log \Delta} \geq (\log \Delta)^{\gamma/2} \\ 2 & \text{otherwise.} \end{cases}$$

Then, the round complexity of Algorithm MWHVC satisfies:

$$T(f, \Delta, \varepsilon) = O\left(f \cdot \log(f/\varepsilon) + \frac{\log \Delta}{\log \log \Delta} + \min\{\log \Delta, f \cdot \log(f/\varepsilon) \cdot (\log \Delta)^\gamma\}\right).$$

Proof. We prove the theorem using a case analysis.

Case $\alpha = \frac{\log \Delta}{f \cdot \log(f/\varepsilon) \cdot \log \log \Delta} \geq 2$ and $\alpha \geq (\log \Delta)^{\gamma/2}$.

This means that the runtime is bounded by $O(\log_\alpha \Delta + f \cdot \log(f/\varepsilon) \cdot \alpha) = O\left(\frac{\log \Delta}{\gamma \cdot \log \log \Delta}\right)$.

Case $\alpha = 2$ and $2 > \frac{\log \Delta}{f \cdot \log(f/\varepsilon) \cdot \log \log \Delta} \geq (\log \Delta)^{\gamma/2}$.

Notice that in this case $2 > (\log \Delta)^{\gamma/2}$ which implies that Δ is constant. Therefore, the round complexity of our algorithm is bounded by $O(\log \Delta + f \cdot \log(f/\varepsilon)) = O(f \cdot \log(f/\varepsilon))$.

Case $\frac{\log \Delta}{f \cdot \log(f/\varepsilon) \cdot \log \log \Delta} < (\log \Delta)^{\gamma/2}$.

This implies that

$$\begin{aligned} \log \Delta &\leq \min\left\{\log \Delta, f \cdot \log(f/\varepsilon) \cdot (\log \Delta)^{\gamma/2} \cdot \log \log \Delta\right\} \\ &= O(\min\{\log \Delta, f \cdot \log(f/\varepsilon) \cdot (\log \Delta)^\gamma\}). \end{aligned}$$

Therefore, since $\alpha = 2$ in this case, the runtime is bounded by

$$O(\log \Delta + f \cdot \log(f/\varepsilon)) = O(f \cdot \log(f/\varepsilon) + \min\{\log \Delta, f \cdot \log(f/\varepsilon) \cdot (\log \Delta)^\gamma\}). \quad \blacktriangleleft$$

Let $W \triangleq \max_v w(v) / \min_v w(v)$. By setting $\varepsilon = 1/(nW)$, we conclude the following result for an f -approximation (recall that the vertex degrees and weights are polynomial in n):

► **Corollary 10.** Algorithm MWHVC computes an f -approximation in $O(f \log n)$ rounds.

Additionally, we get the following range of parameters for which the round complexity is optimal:

► **Corollary 11.** Let $f = O((\log \Delta)^{0.99})$ and $\varepsilon = (\log \Delta)^{-O(1)}$. Then, Algorithm MWHVC computes an $(f + \varepsilon)$ -approximation in $O\left(\frac{\log \Delta}{\log \log \Delta}\right)$ rounds.

For $f = O(1)$ we also get an extension of range of parameters for which the round complexity is optimal. This extension is almost exponential compared to the $\varepsilon = (\log \Delta)^{-O(1)}$ of [5].

► **Corollary 12.** Let $f = O(1)$ and $\varepsilon = 2^{-O((\log \Delta)^{0.99})}$. Then our algorithm computes an $(f + \varepsilon)$ -approximation and terminates in $O\left(\frac{\log \Delta}{\log \log \Delta}\right)$ rounds.

⁵ The statement of the theorem is asymptotic (in Δ). This means that for every constant γ , it holds that either $\log^{\gamma/2} \Delta \geq \log \log \Delta$ or Δ is bounded by a constant (determined by γ) in which case expressions involving Δ can be omitted from the asymptotic expression.

5 Approximation of Covering ILPs

In this section, we present a reduction from solving covering integer linear programs (ILPs) to MWHVC. This reduction implies that one can distributively compute approximate solutions to covering ILPs using a distributed algorithm for MWHVC.

Notation. Let \mathbb{N} denote the set of natural numbers. Let A denote a real $m \times n$ matrix, $\vec{b} \in \mathbb{R}^m$, and $\vec{w} \in \mathbb{R}^n$. Let $LP(A, \vec{b}, \vec{w})$ denote the linear program $\min \vec{w}^T \cdot \vec{x}$ subject to $A \cdot \vec{x} \geq \vec{b}$ and $\vec{x} \geq 0$. Let $ILP(A, \vec{b}, \vec{w})$ denote the integer linear program $\min \vec{w}^T \cdot \vec{x}$ subject to $A \cdot \vec{x} \geq \vec{b}$ and $\vec{x} \in \mathbb{N}^n$.

► **Definition 13.** *The linear program $LP(A, \vec{b}, \vec{w})$ and integer linear program $ILP(A, \vec{b}, \vec{w})$ are covering programs if all the components in A, \vec{b}, \vec{w} are non-negative.*

5.1 Distributed Setting

We denote the number of rows of the matrix A by m and the number of columns by n . Let $f(A)$ (resp., $\Delta(A)$) denote the maximum number of nonzero entries in a row (resp., column) of A .

Given a covering ILP, $ILP(A, \vec{b}, \vec{w})$, the communication network $N(ILP)$ over which the ILP is solved is the bipartite graph $N = (X \times C, E)$, where: $X = \{x_j\}_{j \in [n]}$, $C = \{c_i\}_{i \in [m]}$, and $E = \{(x_j, c_i) \mid A_{i,j} \neq 0\}$. We refer to the nodes in X as variable nodes, and to those in C as constraint nodes. Note that the maximum degree of a constraint node is $f(A)$ and the maximum degree of a variable node is $\Delta(A)$.

We assume that the local input of every variable node x_j consists of w_j and the j 'th column of A (i.e., $A_{i,j}$, for $i \in [m]$). Every constraint vertex c_i is given the value of b_i as its local input. We assume that these inputs can be represented by $O(\log(nm))$ bits. In $(f + 1)$ rounds, every variable node v_j can learn all the components in the rows i of A such that $A_{i,j} \neq 0$ as well as the component b_i .

5.2 Zero-One Covering Programs

The special case in which a variable may be assigned only the values 0 or 1 is called a *zero-one program*. We denote the zero-one covering ILP induced by a matrix A and vectors \vec{b} and \vec{w} by $ZO(A, \vec{b}, \vec{w})$. Every instance of the MWHVC problem is a zero-one program in which the matrix A is the incidence matrix of the hypergraph. The following lemma deals with the converse reduction.

► **Lemma 14.** *Every feasible zero-one covering program $ZO(A, \vec{b}, \vec{w})$ can be reduced to an MWHVC instance with rank $f' < f(A)$ and degree $\Delta' < 2^{f(A)} \cdot \Delta(A)$.*

Proof. Let A_i denote the i 'th row of the matrix A . For a subset $S \subseteq [n]$, let I_S denote the indicator vector of S . Let $\vec{x} \in \{0, 1\}^n$ and let $\sigma_i \triangleq \{j \in [n] \mid A_{i,j} \neq 0\}$. Feasibility implies that $A_i \cdot I_{\sigma_i} \geq b_i$, for every row i . Let \mathcal{S}_i denote the set of all subsets $S \subset [n]$ such that the indicator vector I_S does not satisfy the i 'th constraint, i.e., $A_i \cdot I_S < b_i$. The i 'th constraint is not satisfied, i.e., $A_i \cdot \vec{x} < b_i$, if and only if there exists a set $S \in \mathcal{S}_i$ such that $\vec{x} = I_S$. Hence, $A_i \cdot \vec{x} < b_i$ if and only if the truth value of the following DNF formula is false: $\varphi_i(x) \triangleq \bigvee_{S \in \mathcal{S}_i} \bigwedge_{j \in \sigma_i \setminus S} \text{not}(x_j)$. By De Morgan's law, we obtain that $\text{not}(\varphi_i(x))$ is equivalent to a monotone CNF formula $\psi_i(x)$ such that $\psi_i(x)$ has less than $2^{f(A)}$ clauses, each of which has length less than $f(A)$. We now construct the hypergraph H for the MWHVC instance as follows. For every row i and every $S \in \mathcal{S}_i$, add the hyperedge $e_{i,S} = \sigma_i \setminus S$. (Feasibility implies that $e_{i,S}$ is not empty.) Given a vertex cover C of the

hypergraph, every hyperedge is stabbed by C , and hence I_C satisfies all the formulae $\psi_i(x)$, where $i \in [m]$. Hence, $A_i \cdot I_C \geq b_i$, for every i . The converse direction holds as well, and the lemma follows. \blacktriangleleft

How does $N(ILP)$ (a bipartite graph with $m + n$ vertices) simulate the execution of MWHVC over the hypergraph H ? Each variable node x_j simulates all hyperedges $e_{i,S}$, where $j \in \sigma_i$ and $S \in \mathcal{S}_i$. First, the variable nodes exchange their weights with the variables nodes they share a constraint with in $O(f(A))$ rounds – first every x_j broadcasts its own weight and then each c_i sends all neighbors the weights it received. At each iteration, the variable node sends a raise/stuck message and whether its level was incremented.⁶ Notice that the number of rounds required for each such iteration is $O(1 + f(A)/\log n)$ (i.e., constant for $f(A) = O(\log n)$). Each edge node c_i then broadcasts to all vertices two $f(A)$ -bit messages that indicate two subsets of vertices of the edge: those that sent a raise message (the complement sent a stuck message) and those that incremented their level. Each variable node x_j knows how to update its bid with every $e_{i,S}$ for which $j \in \sigma_i$ and $S \in \mathcal{S}_i$.

We summarize the complexity of the distributed algorithm for solving a zero-one covering program.

\triangleright **Claim 15.** Denoting by $T(f, \Delta, \varepsilon)$ is the running time of Algorithm MWHVC, There exists a distributed CONGEST algorithm for computing an $(f + \varepsilon)$ -approximate solution for zero-one covering programs with running time of $O((1 + f(A)/\log n) \cdot T(f(A), 2^{f(A)} \cdot \Delta(A), \varepsilon))$.

5.3 Reduction of Covering ILPs to Zero-One Covering

Consider the covering ILP $ILP(A, \vec{b}, \vec{w})$. We present a reduction of the ILP to a zero-one covering program.

\blacktriangleright **Definition 16.** Define $M(A, \vec{b}) \triangleq \max_j \max_i \left\{ \frac{b_i}{A_{i,j}} \mid A_{i,j} \neq 0 \right\}$.

We abbreviate and write M for $M(A, \vec{b})$ when the context is clear.

\blacktriangleright **Proposition 17.** Limiting \vec{x} to the box $[0, M]^n$ does not increase the optimal value of the ILP.

\triangleright **Claim 18.** Every covering ILP $ILP(A, \vec{b}, \vec{w})$ can be solved by a zero-one covering program $ZO(A', \vec{b}, \vec{w}')$, where $f(A') \leq f(A) \cdot \lceil \log_2(M) + 1 \rceil$ and $\Delta(A') = \Delta(A)$.

Proof. Let $B = \lceil \log_2 M \rceil$. Limiting each variable x_j by M means that we can replace x_j by B zero-one variables $\{x_{j,\ell}\}_{\ell=0}^{B-1}$ that correspond to the binary representation of x_j , i.e., $x_j = \sum_{\ell=0}^{B-1} 2^\ell \cdot x_{j,\ell}$. This replacement means that the dimensions of the matrix A' are $m \times n'$, where $n' = n \cdot B$. The j 'th column $A^{(j)}$ of A is replaced by B columns, indexed 0 to $B - 1$, where the ℓ 'th column equals $2^\ell \cdot A^{(j)}$. The vector \vec{w}' is obtained by duplicating and scaling the entries of \vec{w} in the same fashion. \blacktriangleleft

Combining Claims 15 and 18, we obtain the following result.

\blacktriangleright **Theorem 19.** There exists a distributed CONGEST algorithm for computing an $(f + \varepsilon)$ -approximate solution for covering integer linear programs $ILP(A, \vec{b}, \vec{w})$ with running time of $O((1 + f(A)/\log n) \cdot T(f(A) \cdot \log M, 2^{f(A)} \cdot M \cdot \Delta(A), \varepsilon))$, where $T(f, \Delta, \varepsilon)$ is the running time of Algorithm MWHVC.

⁶ One needs to modify the MWHVC algorithm slightly so that, in each iteration, the level of every vertex is increased by at most 1. We defer the details to Appendix C.

Proof. Let f_{HVC}, f_{ZO}, F_{ILP} denote the ranks of the hypergraph vertex cover instance, zero-one covering program, and ILP, respectively. We use the same notation for maximum degrees. The reduction of zero-one programs to MWHVC in Lemma 14 implies that $f_{HVC} \leq f_{ZO}$ and $\Delta_{HVC} \leq 2^{f_{ZO}} \cdot \Delta_{ZO}$. The reduction of covering ILPs to zero-one programs in Claim 18 implies that $f_{ZO} \leq f_{ILP} \cdot (1 + \log M)$ and $\Delta_{ZO} \leq \Delta_{ILP}$. The composition of the reductions gives $f_{HVC} \leq f_{ILP} \cdot (1 + \log M)$ and $\Delta_{HVC} \leq 2^{f_{ILP} \cdot (1 + \log M)} \cdot \Delta_{ILP} = 2^{F_{ILP}} \cdot 2M \cdot \Delta_{ILP}$ ◀

After some simplifications, the running time of the resulting algorithm for $(f + \varepsilon)$ -approximate integer covering linear programs is

$$O\left((1 + f/\log n) \cdot \left(\frac{\log \Delta}{\log \log \Delta} + (f \cdot \log M)^{1.01} \cdot \log \varepsilon^{-1} \cdot (\log \Delta)^{0.01}\right)\right).$$

References

- 1 Matti Astrand, Patrik Floréen, Valentin Polishchuk, Joel Rybicki, Jukka Suomela, and Jara Uitto. A Local 2-Approximation Algorithm for the Vertex Cover Problem. In *DISC*, 2009.
- 2 Matti Astrand and Jukka Suomela. Fast distributed approximation algorithms for vertex cover and set cover in anonymous networks. In *SPAA*, 2010.
- 3 Reuven Bar-Yehuda, Keren Censor-Hillel, Mohsen Ghaffari, and Gregory Schwartzman. Distributed Approximation of Maximum Independent Set and Maximum Matching. In *PODC*. ACM, 2017.
- 4 Reuven Bar-Yehuda, Keren Censor-Hillel, and Gregory Schwartzman. A Distributed $(2 + \varepsilon)$ -Approximation for Vertex Cover in $O(\log \Delta / \varepsilon \log \log \Delta)$ Rounds. *J. ACM*, 2017.
- 5 R. Ben-Basat, G. Even, K.-i. Kawarabayashi, and G. Schwartzman. A Deterministic Distributed 2-Approximation for Weighted Vertex Cover in $O(\log n \log \Delta / \log^2 \log \Delta)$ Rounds. In *SIROCCO*, 2018.
- 6 R. Ben-Basat, K.-i. Kawarabayashi, and G. Schwartzman. Distributed Set Cover Approximation: Primal-Dual with Optimal Locality. In *DISC*, 2019.
- 7 Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. An Exponential Separation between Randomized and Deterministic Complexity in the LOCAL Model. In *FOCS*, 2016.
- 8 Richard Cole and Uzi Vishkin. Deterministic Coin Tossing with Applications to Optimal Parallel List Ranking. *Information and Control*, 1986.
- 9 Guy Even, Mohsen Ghaffari, and Moti Medina. Distributed Set Cover Approximation: Primal-Dual with Optimal Locality. In *DISC*, 2018.
- 10 Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In *SODA*. Society for Industrial and Applied Mathematics, 2016.
- 11 Mohsen Ghaffari and Hsin-Hao Su. Distributed Degree Splitting, Edge Coloring, and Orientations. In *SODA*, 2017.
- 12 Fabrizio Grandoni, Jochen Könnemann, and Alessandro Panconesi. Distributed weighted vertex cover via maximal matchings. *ACM Transactions on Algorithms*, 2008. doi:10.1145/1435375.1435381.
- 13 Dorit S. Hochbaum. Approximation Algorithms for the Set Covering and Vertex Cover Problems. *SIAM J. Comput.*, 1982.
- 14 Richard M. Karp. Reducibility Among Combinatorial Problems. In *Proceedings of a symposium on the Complexity of Computer Computations*, 1972.
- 15 Samir Khuller, Uzi Vishkin, and Neal E. Young. A Primal-Dual Parallel Approximation Technique Applied to Weighted Set and Vertex Covers. *J. Algorithms*, 1994.
- 16 Christos Koufogiannakis and Neal E. Young. Distributed algorithms for covering, packing and maximum weighted matching. *Distributed Computing*, 2011.
- 17 Fabian Kuhn. *The price of locality: exploring the complexity of distributed coordination primitives*. PhD thesis, ETH Zurich, 2005.

5:14 Optimal Distributed Covering Algorithms

- 18 Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. The price of being near-sighted. In *SODA*, 2006.
- 19 Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local Computation: Lower and Upper Bounds. *J. ACM*, 2016.
- 20 Alessandro Panconesi and Romeo Rizzi. Some simple distributed algorithms for sparse networks. *Distributed Computing*, 2001.
- 21 Valentin Polishchuk and Jukka Suomela. A simple local 3-approximation algorithm for vertex cover. *Inf. Process. Lett.*, 2009.

A Primal-Dual Approach

The fractional LP relaxation of MWHVC is defined as follows.

$$\begin{aligned}
 & \text{minimize: } \sum_{v \in V} w(v) \cdot x(v) \\
 & \text{subject to:} \\
 & \sum_{v \in e} x(v) \geq 1, \quad \forall e \in E; \quad x(v) \geq 0, \quad \forall v \in V
 \end{aligned} \tag{P}$$

The dual LP is an *Edge Packing* problem defined as follows:

$$\begin{aligned}
 & \text{maximize: } \sum_{e \in E} \delta(e) \\
 & \text{subject to:} \\
 & \sum_{e \ni v} \delta(e) \leq w(v), \quad \forall v \in V; \quad \delta(e) \geq 0, \quad \forall e \in E
 \end{aligned} \tag{D}$$

The following claim is used for proving the approximation ratio of the MWHVC algorithm.

▷ **Claim 20.** Let opt denote the value of an optimal fractional solution of the primal LP (P). Let $\{\delta(e)\}_{e \in E}$ denote a feasible solution of the dual LP (D). Let $\varepsilon \in (0, 1)$ and $\beta \triangleq \varepsilon/(f + \varepsilon)$. Define the β -tight vertices by $T_\varepsilon \triangleq \{v \in V \mid \sum_{e \ni v} \delta(e) \geq (1 - \beta) \cdot w(v)\}$. Then $w(T_\varepsilon) \leq (f + \varepsilon) \cdot \text{opt}$.

Proof.

$$w(T_\varepsilon) = \sum_{v \in T_\varepsilon} w(v) \leq \frac{1}{1 - \beta} \cdot \left(\sum_{v \in T_\varepsilon} \sum_{e \ni v} \delta(e) \right) \leq \frac{f}{1 - \beta} \sum_{e \in E} \delta(e) \leq (f + \varepsilon) \cdot \text{opt}.$$

The last transition follows from $f/(1 - \beta) = f + \varepsilon$ and by weak duality. The claim follows. \triangleleft

B Adaptation to the CONGEST model

We need to show that the message lengths in Algorithm MWHVC are $O(\log n)$.

1. In round 0, every vertex v sends its weight $w(v)$ and degree $|E(v)|$ to every hyperedge in $e \in E(v)$. We assume that the weights and degrees are polynomial in n , hence the length of the binary representations of $w(v)$ and $|E(v)|$ is $O(\log n)$.

Every hyperedge e sends back to every $v \in e$ the pair $(w(v_e), |E(v_e)|)$, where v_e has the smallest normalized weight, i.e., $v_e = \operatorname{argmin}_{v \in e} \{w(v)/|E(v)|\}$.

Every vertex $v \in e$ locally computes $\text{bid}_0(e) = w(v_e)/(2 \cdot |E(v_e)|)$ and $\delta_0(e) = \text{bid}_0(e)$.

2. In round $i \geq 1$, every vertex sends messages. Messages of the sort: “ e is covered”, “raise”, or “stuck” require only a constant number of bits. The increment of v ’s level needs to be sent to the edges in $E(v)$. These increments require $O(\log z) = O(\log n)$ bits.
3. Every edge e sends to every $v \in e$ the number of times that $\text{bid}(e)$ is halved in this iteration. This message is $O(\log z)$ bits long.
4. Every edge sends the final bid to the vertices. Instead of sending the value of the bid, the edge can send a single bit indicating whether the bid was multiplied by α .
5. Finally, if $\alpha = \alpha(e)$ is set locally based on the local maximum degree $\max_{v \in e} |E(v)|$, then every vertex v sends its degree to all the edges $e \in E(v)$. The local maximum degree for e is sent to every vertex $v \in V$, and this parameter is used to compute $\alpha(e)$ locally.

C Algorithm with At Most One Level Increment per Iteration

We propose to do a single change that will ensure that no vertex levels up more than once per iteration. To that end, we modify Line 3f of our MWHVC algorithm to

- For every uncovered edge e , if *all* incoming messages are “raise” $\text{bid}(e) \leftarrow \alpha \cdot \text{bid}(e)$. Send $\text{bid}(e)$ to every $v \in e$, who updates $\delta(e) \leftarrow \delta(e) + \text{bid}(e)/2$.

That is, the algorithm remains intact except that the dual variables $\delta(e)$ are raised by $\text{bid}(e)/2$ rather than by $\text{bid}(e)$. Intuitively, this guarantees that a vertex’s slack does not reduce by more than 50% in each iteration and therefore its level may increase by at most one. We revisit the proof of 2, and, specifically, the *Proof of feasibility of $\delta_{i-1}(e)$ for $i > 1$* .

Proof of feasibility of $\delta_{i-1}(e)$ for $i > 1$. Consider a vertex v . If v joins C in iteration $i - 1$, then $\delta_{i-1} = \delta_{i-2}$, and the packing constraint of v holds by the induction hypothesis. If $v \notin C$, then by $\delta_{i-1}(e) = \delta_{i-2}(e) + \text{bid}_{i-1}(e)/2$, Claim 1, and the induction hypothesis for Eq. 1, we have

$$\begin{aligned} \sum_{e \in E(v)} \delta_{i-1}(e) &= \sum_{e \in E(v)} (\delta_{i-2}(e) + \text{bid}_{i-1}(v)/2) \\ &\leq \left(1 - 0.5^{\ell_{i-1}(v)+1} + 0.5^{\ell_{i-1}(v)+2}\right) \cdot w(v) = \left(1 - 0.5^{(\ell_{i-1}(v)+1)+1}\right) \cdot w(v). \end{aligned} \quad (2)$$

As evident by Eq. 2, the vertex v ’s level may increase by at most once in each iteration.

► **Corollary 21.** *For any iteration $i \geq 1$ and vertex v : $\ell_i(v) \leq \ell_{i-1}(v) + 1$.*

We note that the change to the algorithm does not affect the correctness of claims 1, 2, 4 and Lemma 6. Lemma 7 changes slightly, as there can now be twice as many v -stuck iterations:

► **Lemma 22.** *For every vertex v and level $\ell(v)$, the number of v -stuck iterations is bounded by 2α .*

Proof. Notice that when v reached the level $\ell(v)$, we had $\sum_{e \in E(v)} \delta(e) \geq w(v)(1 - 0.5^{\ell(v)})$. The number of v -stuck iterations is then bounded by the number of times it can send a “stuck” message without reaching $\sum_{e \in E(v)} \delta(e) > w(v)(1 - 0.5^{\ell(v)+1})$. Indeed, once this inequality holds, the level of v is incremented. Every stuck iteration implies, by Line 3e, that $\sum_{e \in E(v)} \text{bid}(e) > \frac{1}{\alpha} \cdot 0.5^{\ell(v)+1} \cdot w(v)$. Therefore, we can bound the number of iteration by $\frac{w(v)(1 - 0.5^{\ell(v)+1}) - w(v)(1 - 0.5^{\ell(v)})}{\frac{1}{\alpha} \cdot 0.5^{\ell(v)+1} \cdot w(v)/2} = 2\alpha$. ◀

We conclude that the algorithm remains correct and its asymptotic complexity does not change.