



A Deterministic Distributed 2-Approximation for Weighted Vertex Cover in $O(\log N \log \Delta / \log^2 \log \Delta)$ Rounds

Ran Ben-Basat^{1(✉)}, Guy Even², Ken-ichi Kawarabayashi³,
and Gregory Schwartzman³

¹ Department of Computer Science, Technion, Haifa, Israel
sran@cs.technion.ac.il

² School of Electrical Engineering, Tel Aviv University, Tel Aviv, Israel
guy@eng.tau.ac.il

³ NII, Tokyo, Japan
{k_keniti,greg}@nii.ac.jp

Abstract. We present a deterministic distributed 2-approximation algorithm for the Minimum Weight Vertex Cover problem in the CONGEST model whose round complexity is $O(\log n \log \Delta / \log^2 \log \Delta)$. This improves over the currently best known deterministic 2-approximation implied by [KVY94]. Our solution generalizes the $(2 + \epsilon)$ -approximation algorithm of [BCS17], improving the dependency on ϵ^{-1} from linear to logarithmic. In addition, for every $\epsilon = (\log \Delta)^{-c}$, where $c \geq 1$ is a constant, our algorithm computes a $(2 + \epsilon)$ -approximation in $O(\log \Delta / \log \log \Delta)$ rounds (which is asymptotically optimal).

1 Introduction

The Minimum Weight Vertex Cover Problem (MWVC) is defined as follows. The input is a graph $G = (V, E)$ with nonnegative vertex weights $w(v)$. A subset $U \subseteq V$ is a *vertex cover* if, for every edge $e = \{u, v\}$, the intersection $U \cap \{u, v\}$ is not empty. The weight of a subset of vertices U is $\sum_{v \in U} w(v)$. The goal is to find a minimum weight vertex cover. This problem is one of the classical NP-hard problems [Kar72].

In this paper we deal with distributed deterministic approximation algorithms for MWVC. We focus on the CONGEST model of distributed computation in which the communication network is the graph G itself.¹ Computation

R. Ben-Basat—This work was partially sponsored by the Technion-HPI research school.

K. Kawarabayashi and G. Schwartzman—This work was supported by JST ERATO Grant Number JPMJER1201, Japan.

¹ In the CONGEST model vertices have distinct IDs (that are polynomial in $|V|$), however, as in [BCS17], our algorithm works also in the case of anonymous vertices.

proceeds in synchronous rounds. Each round consists of three parts: each vertex receives messages from its neighbors, performs a local computation, and sends messages to its neighbors. The sent messages arrive at their destination in the beginning of the next round. In the CONGEST model, message lengths are bounded by $O(\log |V|)$. In order to send vertex weights, we assume that all the vertex weights are positive integers bounded by polynomial in $n \triangleq |V|$. See [BCS17, ĀS10] for detailed overviews of distributed algorithms for MWVC.

Let Δ denote the maximum vertex degree in G . Two of the most relevant results in this setting to our paper are the lower bound of [KMW16] and the upper bound of [BCS17]. The lower bound of Kuhn *et al.* [KMW16] states that every constant approximation algorithm for MWVC requires at least $\Omega(\log \Delta / \log \log \Delta)$ rounds of communication. The upper bound of Bar-Yehuda *et al.* [BCS17] presents a deterministic distributed $(2 + \epsilon)$ -approximation algorithm (BCS Algorithm) that requires $O(\log \Delta / (\epsilon \cdot \log \log \Delta))$ rounds for $\epsilon \in (0, 1)$. For $\epsilon = \Omega(\log \log \Delta / \log \Delta)$, the running time is $O(\log \Delta / \log \log \Delta)$, with no dependence on ϵ , and is optimal according to [KMW16].

In this paper, we present a generalization of the BCS Algorithm with improved guarantees on the running time for certain ranges of ϵ . We focus on decreasing the dependency of the number of rounds on ϵ . Since the round complexity of the BCS Algorithm is optimal for constant values of ϵ (and even $\epsilon = \Omega(\log \log \Delta / \log \Delta)$), we consider values of ϵ that depend on Δ .

Our main result² is a deterministic distributed $(2 + \epsilon)$ -approximation algorithm in which the number of rounds is bounded by

$$O\left(\frac{\log \Delta}{\log \log \Delta} + \frac{\log \epsilon^{-1} \log \Delta}{\log^2 \log \Delta}\right).$$

This result assumes that all the vertices know Δ or an estimate that is a polynomial in Δ . This result leads to the following consequences:

1. If $\epsilon^{-1} = (\log \Delta)^c$, for a constant $c > 0$, then the number of rounds asymptotically matches the lower bound, and is thus optimal. In [BCS17] the same asymptotic running time is guaranteed only for $\epsilon^{-1} = O(\log \Delta / \log \log \Delta)$.
2. If $\epsilon^{-1} = (\log \Delta)^{\omega(1)}$, then the dependency of the round complexity on $1/\epsilon$ is reduced from linear to logarithmic. In addition, the round complexity is decreased by an additional factor of $\log \log \Delta$.
3. Every $(2 + \epsilon)$ -approximation is a 2-approximation if $\epsilon < 1/(nW)$, where $W = \max_v w(v)$. Since we assume that $W = n^{O(1)}$, where $n = |V|$, we obtain a 2-approximate deterministic distributed algorithm for MWVC with round complexity $O(\log n \cdot \log \Delta / \log^2 \log \Delta)$. This improves over the 2-approximation in $O(\log^2 n)$ rounds implied by [KVY94]³ (which has the lowest round complexity for deterministic 2-approximation to the best of our knowledge).

² All logarithms are base 2 unless the basis is written explicitly.

³ The actual result is stated as a $(2 + \epsilon)$ -approximation in $O(\log \epsilon^{-1} \log n)$ rounds, from which we infer a 2-approximation by setting $\epsilon = 1/nW$.

Our round complexity increases for the case that the maximum degree Δ is unknown to the vertices of the graphs. We propose two alternatives for the case that Δ is unknown. The first alternative holds for every $\epsilon \in (0, 1)$, and achieves a $(2 + \epsilon)$ -approximate solution with a round complexity of $O\left(\frac{\log \epsilon^{-1} \log \Delta}{\log \log \Delta}\right)$. The second alternative holds for $\epsilon > (\log \Delta)^q$, where $q > 0$ is a constant. In the second alternative, a $(2 + \epsilon)$ -approximation is achieved with an optimal asymptotic round complexity of $O(\log \Delta / \log \log \Delta)$.

Our algorithm builds on the BCS Algorithm [BCS17]. This algorithm adapts the *local ratio* framework [BE85] to the distributed setting, with several improvements that provide the desired speedup. The BCS Algorithm can be also interpreted as the following “primal-dual” algorithm. Essentially the algorithm aims to increase the edge variables (i.e., *dual*) such that the following holds:

1. The sum of edge variables incident to every node does not exceed its weight (*feasibility of edge variables*).
2. The set of vertices whose edge variable sum is at least $(1 - \epsilon)$ -fraction of the vertex weight constitute a *vertex cover*.

The above conditions yield a $(2 + \epsilon)$ -approximation for MWVC.

The challenge in the above framework is to maintain feasibility of the edge variables while converging as fast as possible to a vertex cover. To increase the edge variables, vertices send *offers* to their neighbors. The neighbors respond to these offers in a way that guarantees feasibility of the edge variables. This requires a coordination mechanism in the distributed setting, as a vertex both sends and receives offers simultaneously. To this end, the weight of every vertex is divided into two parts: *vault* and *bank*. Offers are allocated from the vault, while responses are allocated from the bank, respectively. Hence the agreed upon increases to the edge variables do not violate the feasibility of the edge variables. The BCS algorithm sets the vault to be an ϵ -fraction of the vertex weight (and the bank to be the remainder). This leads to a running time of $O(\epsilon^{-1} \log \Delta / \log \log \Delta)$ and $O(\log \Delta / \log \log \Delta)$ if $\epsilon = \Omega(\log \log \Delta / \log \Delta)$.

Our algorithm introduces three modifications to the BCS Algorithm, which allows us to improve the round complexity. First, we attach *levels* to the vertices that measure by how much the remaining weight of a vertex has decreased. Second, the size of the vault decreases as the level of the vertex increases. Third, offers are not sent to all the neighbors. Instead, offers are sent only to the neighbors whose level is the smallest level among the remaining neighbors.

Related Work. An excellent overview of the related work is presented in [BCS17, ÅS10] which we summarize hereinafter. Minimum vertex cover is one of Karp’s 21 NP-hard problems [Kar72]. A simple 2-approximation for the unweighted version can be achieved by a reduction from maximal matching (see, e.g., [CLRS09, GJ79]). For the weighted case, [BE81] achieves the first linear-time 2-approximation algorithm using the primal-dual schema, while [BE85] achieves the same result using the local-ratio technique. Prior to that, the first polynomial-time 2-approximation algorithm was due to [NJ75] and observed

by [Hoc82]. For any constant $\epsilon > 0$, if the Unique Games conjecture holds, no polynomial-time algorithm can compute a $(2 - \epsilon)$ approximation of the minimum vertex cover [KR08].

Let us now turn our attention to the distributed setting. Let us start from the unweighted case. A 2-approximation can be found in $O(\log^4 n)$ rounds [HKP01] and in $O(\Delta + \log^* n)$ rounds [PR01]. Completely local algorithms with no dependence on n are presented in [AFP+09] which gives an $O(\Delta^2)$ -round 2-approximation algorithm, and in [PS09] which gives an $O(\Delta)$ -round 3-approximation algorithm. Using the maximal matching algorithm of [BEPS12] gives a 2-approximation algorithm for vertex cover in $O(\log \Delta + (\log \log n)^4)$ rounds. This can be made into a $(2 + 1/\text{poly} \Delta)$ -approximation algorithm within $O(\log \Delta)$ rounds [Pet16].

For the weighted case, [GKP08] presents a randomized 2-approximation algorithm in $O(\log n + \log W)$ rounds (where W is a bound on the vertex weights). In [KY11] the first (randomized) 2-approximation algorithm running in $O(\log n)$ rounds is presented (note that the running time is logarithmic in n and independent of the weights). A deterministic 2-approximation algorithm in $O(\Delta + \log^* n)$ rounds is given within [PR01]. In [KVY94], a deterministic $(2 + \epsilon)$ -approximation algorithm is given within $O(\log \epsilon^{-1} \log n)$ rounds. As for deterministic algorithms independent of n , [KMW06] presents a $(2 + \epsilon)$ -approximation algorithm in $O(\epsilon^{-4} \log \Delta)$ rounds and [AFP+09] presents a 2-approximation algorithm in $O(1)$ rounds for $\Delta \leq 3$, while [AS10] presents a 2-approximation algorithm in $O(\Delta + \log^* W)$ rounds (where $W \triangleq \max_v w(v)$). Finally in [BCS17] a deterministic $(2 + \epsilon)$ -approximation which runs in $O(\epsilon \log \Delta / \log \log \Delta)$ rounds is given. In [Sol18] a $(2 + \epsilon)$ -approximation in $O(\epsilon^{-1} \log(\alpha/\epsilon) / \log \log(\alpha/\epsilon))$ rounds for graphs of arboricity bounded by α .

As the result of [Sol18] uses the algorithm of [BCS17] as a black box, plugging $\Delta = \alpha/\epsilon$, our results can also be used. This means all of results stated in this paper also hold for bounded arboricity graphs setting $\Delta = \alpha/\epsilon$. We list the previous results and the results of this paper in Table 1 (Adapted from [AS10]).

2 The MWVC Local Ratio Template

In this section we overview [BCS17]'s local ratio paradigm for approximating MWVC. We note that the template does not assume anything about the model of computation and that our algorithms will fit into this framework. This template can also be viewed via the primal-dual schema.

Let $G = (V, E)$ denote a graph with a vertex-weight function $w : V \rightarrow \mathbb{R}^+$. An edge-weight function $\delta : E \rightarrow \mathbb{R}^+$ is G -valid if for every vertex v the incident edges weight sum does not exceed $w(v)$; that is, δ is G -valid if $\forall v \in V : \sum_{v \ni e} \delta(e) \leq w(v)$. (In fact, a G -valid function δ is a feasible solution to the dual edge packing LP.)

Next, for a G -valid function δ , define the vertex-weight function $\tilde{w}_\delta : V \rightarrow \mathbb{R}^+$ by $\tilde{w}_\delta(v) = \sum_{e: v \in e} \delta(e)$. Let $S_\delta = \{v \in V \mid w(v) - \tilde{w}_\delta(v) \leq \epsilon' w(v)\}$ be the set of vertices for which w and \tilde{w} differ by at most $\epsilon' w(v)$, for $\epsilon' = \epsilon/(2 + \epsilon)$. We refer

Table 1. In the table (adapted from [ÅS10]), $n = |V|$ and $\epsilon \in (0, 1)$. The running times are stated for the case of unit weight vertices. For randomized algorithms the running times hold in expectation or with high probability.

Deterministic	Weighted	Approximation	Time ($W = 1$)	Algorithm
no	yes	2	$O(\log n)$	[GKP08]
no	yes	2	$O(\log n)$	[KY09]
yes	no	3	$O(\Delta)$	[PS09]
yes	no	2	$O(\log^4 n)$	[HKP01]
yes	no	2	$O(\Delta^2)$	[AFP+09]
yes	yes	$2 + \epsilon$	$O(\log \epsilon^{-1} \log n)$	[KVY94]
yes	yes	2	$O(\log^2 n)$	[KVY94]
yes	yes	$2 + \epsilon$	$O(\epsilon^{-4} \log \Delta)$	[Hoc82, KMW06]
yes	yes	2	$O(\Delta + \log^* n)$	[PR01]
yes	yes	2	$O(\Delta)$	[ÅS10]
yes	yes	2	$O(1)$ for $\Delta \leq 3$	[AFP+09]
yes	yes	$2 + \epsilon$	$O(\epsilon^{-1} \log \Delta / \log \log \Delta)$	[BCS17]
yes	yes	$2 + \frac{\log \log \Delta}{\log \Delta}$	$O(\log \Delta / \log \log \Delta)$	[BCS17]
yes	yes	$2 + \epsilon$	$O\left(\frac{\log \Delta}{\log \log \Delta} + \frac{\log \epsilon^{-1} \log \Delta}{\log^2 \log \Delta}\right)$	This work
yes	yes	$2 + (\log \Delta)^{-c}$	$O(\log \Delta / \log \log \Delta)$	This work, $\forall c = O(1)$
yes	yes	2	$O(\log n \log \Delta / \log^2 \log \Delta)$	This work

to vertices in S_δ as ϵ' -tight vertices. The essence of the template consists of two parts: (1) The sum of the weights of the vertices in S_δ is at most $(2 + \epsilon)$ times the weight of an optimal solution to MWVC. (2) When the algorithm terminates, S_δ is a vertex cover.

Theorem 1. ([BCS17]) *Fix $\epsilon > 0$ and let δ be a G -valid function. Let OPT be the sum of weights of vertices in a minimum weight vertex cover S_{OPT} of G . Then $\sum_{v \in S_\delta} w(v) \leq (2 + \epsilon)OPT$. In particular, if S_δ is a vertex cover, then it is a $(2 + \epsilon)$ -approximation for MWVC for G .*

3 A Fast Distributed Implementation

In this section, we present a modification of the distributed algorithm for MWVC of Bar-Yehuda *et al.* [BCS17]. The pseudo-code for our algorithm is given in Algorithm 1. In this section we assume that the maximal degree Δ is known to all vertices.⁴ In Sect. 4 we provide an algorithm with a slightly higher running time in which this assumption is lifted.

For clarity of presentation, we first describe an implementation for the LOCAL model. This algorithm can be easily adapted to the CONGEST model using the techniques of [BCS17].

⁴ A polynomial upper bound of $\Delta^{O(1)}$ would yield the same asymptotic bound on the number of rounds.

Overview of Algorithm 1. The algorithm uses the following parameters: (i) $\epsilon' \triangleq \epsilon/(2 + \epsilon)$, (ii) $\gamma \in (0, 1)$, (iii) $z \triangleq \lceil \log_\gamma \epsilon' \rceil$. The parameter ϵ' is used for defining tightness of the dual packing constraint. The parameter γ is used for defining levels. Loosely speaking, in every iteration the weight of a vertex is reduced, and the level of a vertex is proportional to $\log_\gamma(w_i(v)/w_0(v))$. The parameter z is used to bound the number of levels till a vertex becomes ϵ' -tight, meaning that $w_i(v)/w_0(v) \leq \epsilon'$.

Our algorithm, listed as Algorithm 1, is a variation of the Algorithm of Bar-Yehuda *et al.* [BCS17] with a few modifications. We begin with a description of the common features. In the course of the algorithm, the weight of each vertex is reduced. Once a vertex v becomes ϵ' -tight (i.e., the reduced weight is an ϵ' -fraction of its original weight) it decides to join the vertex cover and terminates after sending the message (v, cover) to its remaining neighbors. The message (v, cover) causes the neighbors of v to erase v from their list of remaining neighbors. If a vertex v loses all its neighbors (i.e., becomes isolated), it decides that v is not in the vertex cover, and terminates. Upon termination, the ϵ' -tight vertices constitute a vertex cover.

The handling of offers is as in [BCS17]. Vertex v sends an irrevocable offer $\text{request}_i(v, u)$ to every $u \in N_i(v)$. The offers are allocated from the vault. The responses to the offers are allocated greedily from the bank, namely v 's responses satisfy: $\text{budget}_i(v, u) \leq \text{request}_i(u, v)$ and $\sum_u \text{budget}_i(v, u) \leq \text{bank}_i(v)$. The updating of the weights can be interpreted as follows. For every edge $e = \{u, v\}$ the dual edge packing variable $\delta(e)$ is increased by $\text{budget}_i(u, v) + \text{budget}_i(v, u)$. The remaining weight satisfies $w_{i+1}(v) = w_0(v) - \sum_{e \ni v} \delta(e)$. Note that each iteration of the while-loop requires a constant number of communication rounds.

The first modification is that we attach a *level* to each vertex as follows. Let $w_i(v)$ denote the weight of v in the beginning of iteration i of the while-loop. The level of v in iteration i satisfies $\ell_i(v) = 1 + \left\lceil \log_\gamma \frac{w_i(v)}{w_0(v)} \right\rceil$. Note that the initial level is one, and that if the level of v is greater than z , then v is ϵ' -tight (see Claim 3.1).

The second modification is how we partition $w_i(v)$ between the vault and the bank. Instead of using a fixed fraction of the initial weight for the vault, our vault decreases as the level of the vertex increases. Formally, $\text{vault}_i(v) \triangleq w_0(v) \cdot \gamma^{\ell_i(v)}$. The bank is the rest of the weight, namely, $\text{bank}_i(v) \triangleq w_i(v) - \text{vault}_i(v)$.

The third modification is that in each iteration, every vertex v only sends offers to its remaining neighbors with the smallest level. Let $N_i(v)$ denote the set of remaining neighbors of v in the beginning of the i th iteration. The smallest level of the neighbors of v is defined by $\ell'_i(v) \triangleq \min\{\ell_i(u) \mid u \in N_i(v)\}$. The set of neighbors of lowest level is defined by $N'_i(v) \triangleq \{u \in N_i(v) \mid \ell_i(u) = \ell'_i(v)\}$. Let $d'_i(v) = |N'_i(v)|$. The size of each offer sent is $\text{vault}_i(v)/d'_i(v)$.

Note that if $\gamma = \epsilon'$, then Algorithm 1 reduces to the BCS Algorithm because there is just one level, and the vault size is fixed and equals to $\epsilon' \cdot w_0(v)$. On the other hand, if $\gamma = 1/2$, then there are $O(\log 1/\epsilon)$ levels. Per level $\ell_i(v)$, the algorithm can be viewed as a version of the BCS algorithm with $\epsilon' = 1/2^{\ell_i(v)}$. This also explains why our algorithm may be adapted to the CONGEST model

of distributed computation using the techniques of [BCS17]. In essence they give an adaptation for a single level of our algorithm, which can easily be extended to multiple levels.

We now state the main theorem of this work.

Theorem 2. *Algorithm 1 (with $\gamma = \frac{1}{\sqrt{\log \Delta}}$ if $\Delta > 16$ and $\gamma = 0.5$ otherwise) is a deterministic distributed $(2 + \epsilon)$ -approximation algorithm for MWVC. The number of rounds required for the algorithm to terminate is $O\left(\frac{\log \Delta}{\log \log \Delta} + \frac{\log \epsilon^{-1} \log \Delta}{\log^2 \log \Delta}\right)$ if $\Delta > 16$ and $O(\log \epsilon^{-1})$ otherwise.*

3.1 Proof of Theorem 2

Notation. In the analysis we use $w_i(v), \ell_i(v), d'_i(v)$ to denote the value of these variables at the beginning of the i th iteration.

The following claim states an invariant that Algorithm 1 satisfies.

Claim. The following invariant holds in every iteration of the while-loop:

$$\gamma^{\ell_i(v)} < \frac{w_i(v)}{w_0(v)} \leq \gamma^{\ell_i(v)-1} \tag{1}$$

Hence, (i) $vault_i(v) < w_i(v)$ and (ii) if $\ell_{i+1}(v) \geq z + 1$, then $\frac{w_{i+1}(v)}{w_0(v)} \leq \epsilon'$.

This invariant of Claim 3.1 implies, among other things, that every vertex that decides to join the vertex cover is ϵ' -tight. This property, together with the fact that the set of vertices that join the vertex cover constitute a vertex cover leads to the proof that Algorithm 1 is a $(2 + \epsilon)$ -approximation algorithm. An analogous lemma and its proof also appears in [BCS17]. We remark that termination of the algorithm is implied by the upper bound on the number of iterations of the while-loop proved in the sequel.

Lemma 1. ([BCS17, Lemma 3.2]) *For every $\epsilon, \gamma \in (0, 1)$, upon termination Algorithm 1 computes a $(2 + \epsilon)$ -approximate solution to MWVC.*

In the following lemma we show that, for every vertex v and every iteration of the while-loop, either many of v 's neighbors of the smallest level have increased their level or v 's weight has decreased significantly.

Lemma 2. *Let $K > 1$. Let i be an iteration of the while-loop in the execution of Algorithm 1 by vertex v in which v does not join the cover. At least one of following conditions must hold:*

1. *At least $d'_i(v)(1 - 1/K)$ of the neighbors of v of the lowest level have increased their level. Formally, If $\ell'_{i+1}(v) = \ell'_i(v)$, then $d'_{i+1}(v) < d'_i(v)/K$.*
2. *$w_{i+1}(v) \leq w_i(v) - w_0(v)\gamma^{\ell_i(v)}/K$.*

Algorithm 1. A distributed $(2 + \epsilon)$ -approximation algorithm for MWVC, code for vertex v . (Listing taken from [BCS17] and edited to include our modifications.)

```

1   $\gamma$  = parameter in the interval  $(0, 1)$ .
2   $\epsilon' = \epsilon / (2 + \epsilon)$ 
3   $z = \lceil \log_\gamma \epsilon' \rceil$ 
4   $w_0(v) = w(v)$ 
5   $\ell_0(v) = 1$ 
6   $N_0(v) = N(v)$ ,  $d_i(v) \triangleq |N_i(v)|$ 
7  //Let  $N'_i(v)$  be the set of neighbors of lowest level in iteration  $i$  and
    $d'_i(v) \triangleq |N'_i(v)|$ 
8   $i = 0$ 
9  while true do
10 |    $vault_i(v) = w_0(v) \cdot \gamma^{\ell_i(v)}$ 
11 |    $bank_i(v) = w_i(v) - vault_i(v)$ 
12 |    $w_{i+1}(v) = w_i(v)$  and  $\ell_{i+1}(v) \leftarrow \ell_i(v)$ 
13 |   foreach  $u \in N'_i(v)$  do
14 |      $request_i(v, u) = vault_i(v) / d'_i(v)$ 
15 |     Send  $request_i(v, u)$  to  $u$ 
16 |     Let  $budget_i(u, v)$  be the response from  $u$ 
17 |      $w_{i+1}(v) = w_{i+1}(v) - budget_i(u, v)$ 
18 |   Let  $u_1 \dots u_{m_i}$  be an arbitrary order of neighbors that sent requests in this
   iteration
19 |   foreach  $k = 1, \dots, m_i$  do
20 |     Let  $request_i(u_k, v)$  be received from  $u_k \in N'_i(v)$ 
21 |      $budget_i(v, u_k) = \min\{request_i(u_k, v), bank_i(v) - \sum_{t=1}^{k-1} budget_i(v, u_t)\}$ 
22 |     Send  $budget_i(v, u_k)$  to  $u_k$ 
23 |    $w_{i+1}(v) = w_{i+1}(v) - \sum_{k=1}^{d_i(v)} budget_i(v, u_k)$ 
24 |   if  $w_{i+1}(v) \neq 0$  and  $w_{i+1}(v) \leq vault_i(v)$  then
25 |      $\ell_{i+1}(v) = 1 + \lceil \log_\gamma \frac{w_{i+1}(v)}{w_0(v)} \rceil$ 
26 |   if  $w_{i+1}(v) = 0$  or  $\ell_{i+1}(v) \geq z + 1$  then
27 |     Send  $(v, cover)$  to all neighbors
28 |     Return InCover
29 |   foreach  $(u, cover)$  received from  $u \in N_i(v)$  do
30 |      $N_i(v) = N_i(v) \setminus \{u\}$ 
31 |   if  $d_i(v) = 0$  then
32 |     Return NotInCover
33 |    $N_{i+1}(v) = N_i(v)$ 
34 |    $i = i + 1$ 

```


Proof. Assume that $\ell'_{i+1}(v) = \ell'_i(v)$ and $d'_{i+1}(v) \geq d'_i(v)/K$. Note that if the level of a vertex remains unchanged (i.e., $\ell'_{i+1}(v) = \ell'_i(v)$), then either $w_{i+1}(v) = 0$ or $w_{i+1}(v) > vault_i(v)$. If $w_{i+1}(v) = 0$, then v joins the cover, a contradiction. If $w_{i+1}(v) > vault_i(v)$, then the bank was not exhausted and $budget_i(u, v) = request_i(v, u)$. To conclude, at least $d'_i(v)/K$ vertices $u \in N'_i(v)$ responded with $budget_i(u, v) = request_i(v, u)$. This implies that

$$\begin{aligned} w_i(v) - w_{i+1}(v) &\geq \frac{d'_i(v)}{K} \cdot \frac{vault_i(v)}{d'_i(v)} \\ &= w_0(v)\gamma^{\ell_i(v)}/K. \end{aligned}$$

Lemma 3. *For every $\gamma \in (0, 1)$ and $K > 1$, the number of iterations of the while-loop for every vertex v is bounded by:*

$$z \cdot \left(\frac{K}{\gamma} + \frac{\log d(v)}{\log K} \right).$$

Proof. The number of levels is bounded by z . Hence it suffices to prove that the number of rounds per level is at most $K/\gamma + \log_K d(v)$. Indeed, the number of rounds that satisfy Condition 1 per level is bounded by $\log_K d(v)$ because $d'_i(v)$ is divided by at least K in each such iteration.

We now bound the number of iterations that satisfies Condition 2 per level. By Claim 3.1, $w_0(v) \cdot \gamma^{\ell_i(v)-1} \geq w_i(v)$. Hence, the number of iterations that satisfies Condition 2 is bounded by K/γ , as required, and the lemma follows.

We now prove Theorem 2.

Proof. First, consider the case where $\Delta \leq 16$. We set $\gamma = 0.5$ (hence, $z = O(\log \epsilon^{-1})$) and $K = 2$. Lemma 3 immediately shows that the termination time is $O(\log \epsilon^{-1})$. Next, assume that $\Delta > 16$ (thus hereafter: $\log \log \Delta > 2$, $\log \Delta / \log \log \Delta > 2$, $1/\sqrt{\log \Delta} < 1/2$, and $\sqrt{\log \Delta} / \log \log \Delta > 1$). We set $\gamma = 1/\sqrt{\log \Delta}$ and $K = \sqrt{\log \Delta} / \log \log \Delta$. Now we can express the running time as:

$$\begin{aligned} z \cdot \left(\frac{K}{\gamma} + \frac{\log d(v)}{\log K} \right) &\leq z \left(K\gamma^{-1} + \frac{\log \Delta}{\log K} \right) = z \left(\frac{\log \Delta}{\log \log \Delta} + \frac{\log \Delta}{0.5 \log \log \Delta - \log \log \log \Delta} \right) \\ &= O \left(\frac{z \log \Delta}{\log \log \Delta} \right). \end{aligned}$$

Let us analyze the running time according to the values of ϵ . First, consider the case where $\epsilon^{-1} = \log^{O(1)} \Delta$. Since $\epsilon \in (0, 1)$, it follows that $\epsilon' = \Theta(\epsilon)$. We get that $z < 1 + \log_\gamma \epsilon' = O(1 + \log \epsilon^{-1} / \log \gamma^{-1}) = O(1 + \log \log \Delta / \log \log \Delta) = O(1)$. Thus, the total running time for this case is $O(\log \Delta / \log \log \Delta)$. Next we consider the complementary case, where $\epsilon^{-1} = \log^{\omega(1)} \Delta$. This means that $\log \epsilon^{-1} = \omega(\log \log \Delta)$. Therefore, $z = O(\log \epsilon^{-1} / \log \gamma^{-1}) = O(\log \epsilon^{-1} / \log \log \Delta)$, and the running time for the second case is given by $O(\log \epsilon^{-1} \log \Delta / \log^2 \log \Delta)$. Thus, we may express the running time of our algorithm asymptotically as:

$$O(\log \Delta / \log \log \Delta + \log \epsilon^{-1} \log \Delta / \log^2 \log \Delta).$$

The number of rounds is bounded as required, and the theorem follows.

4 An Algorithm Without Knowing Δ

The bound on the round complexity in Theorem 2 assumes that every vertex knows the maximum degree Δ (or a polynomial upper bound on Δ). This is required in order to determine the value of γ . In this section we consider the setting in which Δ is unknown to the vertices.

Note that the analysis in Lemma 3 is per a vertex. Hence, in the analysis of the round complexity, we may use a different K per a vertex. Let K_v denote the value of K that is used in the analysis for bounding the round complexity of v .

We propose two alternatives for the setting of unknown maximum degree, as follows:

1. In the first setting, we simply set $\gamma = 1/2$ in the algorithm. For the analysis, we consider $K_v = \frac{\log d(v)}{\log \log d(v)}$, where $d(v)$ denotes the degree of v . Plugging in these parameters in Lemma 3 gives a round complexity of $O\left(\frac{\log \epsilon^{-1} \log d(v)}{\log \log d(v)}\right)$.
2. For any $q = O(1)$, we can set $\gamma = \epsilon^{1/2q}$ (hence, $z = O(1)$). An analysis with $K_v = \frac{\gamma \log d(v)}{\log \log d(v)}$ shows that v terminates in the optimal $O(\log d(v)/\log \log d(v))$ rounds for any $\epsilon > (\log \Delta)^{-q}$. This is because

$$K_v = \frac{\epsilon^{1/2q} \log d(v)}{\log \log d(v)} > \frac{\log^{-0.5} d(v) \log d(v)}{\log \log d(v)} = \frac{\log^{0.5} d(v)}{\log \log d(v)}.$$

That allows us to express the running time as

$$z(K_v \gamma^{-1} + \log d(v)/\log K_v) = O(\log d(v)/\log \log d(v)).$$

References

- [ÅFP+09] Åstrand, M., Floréen, P., Polishchuk, V., Rybicki, J., Suomela, J., Uitto, J.: A local 2-approximation algorithm for the vertex cover problem. In: Keidar, I. (ed.) DISC 2009. LNCS, vol. 5805, pp. 191–205. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04355-0_21
- [ÅS10] Åstrand, M., Suomela, J.: Fast distributed approximation algorithms for vertex cover and set cover in anonymous networks. In: SPAA 2010 Proceedings of the 22nd Annual ACM Symposium on Parallelism in Algorithms and Architectures, Thira, Santorini, Greece, pp. 294–302, 13–15 June 2010
- [BCS17] Bar-Yehuda, R., Censor-Hillel, K., Schwartzman, G.: A distributed $(2 + \epsilon)$ -approximation for vertex cover in $o(\log \Delta / \epsilon \log \log \Delta)$ rounds. *J. ACM* **64**(3), 23:1–23:11 (2017)
- [BE81] Bar-Yehuda, R., Even, S.: A linear-time approximation algorithm for the weighted vertex cover problem. *J. Algorithms* **2**(2), 198–203 (1981)
- [BE85] Bar-Yehuda, R., Even, S.: A local-ratio theorem for approximating the weighted vertex cover problem. *N.-Holland Math. Stud.* **109**, 27–45 (1985)
- [BEPS12] Barenboim, L., Elkin, M., Pettie, S., Schneider, J.: The locality of distributed symmetry breaking. In: 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, pp. 321–330, 20–23 October 2012

- [CLRS09] Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: *Introduction to Algorithms*, 3rd Edition. The MIT Press, Cambridge (2009)
- [GJ79] Garey, M.R., Johnson, D.S., Freeman, W.H.: *Computers and Intractability: A Guide to the Theory of NP-Completeness* (1979)
- [GKP08] Grandoni, F., Könemann, J., Panconesi, A.: Distributed weighted vertex cover via maximal matchings. *ACM Trans. Algorithms* **5**(1) (2008)
- [HKP01] Hanckowiak, M., Karonski, M., Panconesi, A.: On the distributed complexity of computing maximal matchings. *SIAM J. Discrete Math.* **15**(1), 41–57 (2001)
- [Hoc82] Hochbaum, D.S.: Approximation algorithms for the set covering and vertex cover problems. *SIAM J. Comput.* **11**(3), 555–556 (1982)
- [Kar72] Karp, T.M.: Reducibility among combinatorial problems. In: *Proceedings of a Symposium on the Complexity of Computer Computations*. The IBM Thomas J. Watson Research Center, Yorktown Heights. Springer, New York, pp. 85–103, 20–22 March 1972. https://doi.org/10.1007/978-1-4684-2001-2_9
- [KMW06] Kuhn, F., Moscibroda, T., Wattenhofer, R.: The price of being near-sighted. In: *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA*, pp. 980–989, 22–26 January 2006
- [KMW16] Kuhn, F., Moscibroda, T., Wattenhofer, R.: Local computation: lower and upper bounds. *J. ACM* **63**(2), 17:1–17:44 (2016)
- [KR08] Khot, S., Regev, O.: Vertex cover might be hard to approximate to within 2-epsilon. *J. Comput. Syst. Sci.* **74**(3), 335–349 (2008)
- [KVY94] Khuller, S., Vishkin, U., Young, N.E.: A primal-dual parallel approximation technique applied to weighted set and vertex covers. *J. Algorithms* **17**(2), 280–289 (1994)
- [KY09] Koufogiannakis, C., Young, N.E.: Distributed and parallel algorithms for weighted vertex cover and other covering problems. In: *Proceedings of the 28th ACM Symposium on Principles of Distributed Computing, PODC 2009*, pp. 171–179. ACM, New York (2009)
- [KY11] Koufogiannakis, C., Young, N.E.: Distributed algorithms for covering, packing and maximum weighted matching. *Distrib. Comput.* **24**(1), 45–63 (2011)
- [NJ75] Nemhauser, G.L., Trotter Jr., L.E.: Vertex packings: structural properties and algorithms. *Math. Program.* **8**(1), 232–248 (1975)
- [Pet16] Pettie, S.: Personal communication (2016)
- [PR01] Panconesi, A., Rizzi, R.: Some simple distributed algorithms for sparse networks. *Distrib. Comput.* **14**(2), 97–100 (2001)
- [PS09] Polishchuk, V., Suomela, J.: A simple local 3-approximation algorithm for vertex cover. *Inf. Process. Lett.* **109**(12), 642–645 (2009)
- [Sol18] Solomon, S.: Local algorithms for bounded degree sparsifiers in sparse graphs. In: *ITCS*, volume 94 of LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 52:1–52:19 (2018)