A Massively Parallel Algorithm for Minimum Weight Vertex Cover

Mohsen Ghaffari
ETH Zurich
ghaffari@inf.ethz.ch

Ce Jin
Tsinghua University
jinc16@mails.tsinghua.edu.cn

Daan Nilis
ETH Zurich
nilisdaan@gmail.com

ABSTRACT

We present a massively parallel algorithm, with near-linear memory per machine, that computes a $(2 + \epsilon)$ -approximation of minimum-weight vertex cover in $O(\log \log d)$ rounds, where d is the average degree of the input graph.

Our result fills the key remaining gap in the state-of-the-art MPC algorithms for vertex cover and matching problems; two classic optimization problems, which are duals of each other. Concretely, a recent line of work—by Czumaj et al. [STOC'18], Ghaffari et al. [PODC'18], Assadi et al. [SODA'19], and Gamlath et al. [PODC'19]—provides $O(\log \log n)$ time algorithms for $(1 + \epsilon)$ -approximate maximum weight matching as well as for $(2 + \epsilon)$ -approximate minimum cardinality vertex cover. However, the latter algorithm does not work for the general *weighted* case of vertex cover, for which the best known algorithm remained at $O(\log n)$ time complexity.

CCS CONCEPTS

• **Theory of computation** → **Massively parallel algorithms.**

KEYWORDS

massively parallel computation, weighted vertex cover, approximation algorithms, round compression

ACM Reference Format:

Mohsen Ghaffari, Ce Jin, and Daan Nilis. 2020. A Massively Parallel Algorithm for Minimum Weight Vertex Cover. In *Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '20)*, July 15–17, 2020, Virtual Event, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3350755.3400260>

1 INTRODUCTION

Over the past decade, and sparked by the practical successes of popular computing platforms such as MapReduce [9], Hadoop [25], Dryad [17] and Spark [27], the Massively Parallel Computation (MPC) model has emerged as a theoretical abstraction for large-scale parallel computation and it is receiving increasingly more attention from the algorithmic community. In contrast to the fine-grained parallelism found in celebrated models such as the PRAM [26], this new model allows for a higher-level of granularity. In particular, instead of breaking computation into small read or write operations

(and basic calculations) that are assumed to happen in lock-step rounds across all processors, the model assumes that each machine can process a small (e.g., polynomially smaller than the entire input) chunk of data per time unit and focuses on the number of rounds of parallelism in such a coarse-grained view. We formally introduce the model in Section 1.1.

Shortly after its inception, it has been shown that MPC is at least as powerful as PRAM [15, 19]. Soon after, significantly faster algorithms were developed for a number of important graph problems. The exact speedup depends on the memory that one machine has with respect to the total number of vertices present in the graph. In the setting where each machine has roughly a linear amount of memory with respect to the number of vertices, many of the central problems have been solved in $O(\log \log n)$ rounds. This is true for problems including Maximal Independent Set, nearly Maximum Matching, Maximal Matching, and 2-approximate Minimum Cardinality Vertex Cover [7, 8, 13]. The general technique is a certain “round compression”, which starts with an $O(\log n)$ round PRAM or LOCAL [22] model algorithm and successively compresses a considerable fraction of the remaining rounds of the algorithm into just one round of the MPC model. However, these results do not extend to the Minimum Weight Vertex Cover problem and particularly the compression technique fails in managing the deviations created in the weighted case. Indeed, prior to this work, the question of a similar algorithm for the weighted case remained open, and the best known algorithm for the weighted case remained at the classic $O(\log n)$ time that follows from PRAM and LOCAL model literature (e.g. [20]). In this paper, we resolve this problem by presenting an $O(\log \log n)$ time algorithm for the weighted case.

In the next sections we formally introduce the MPC model and the current state-of-the-art. Consequently we describe our technical contributions on a high level.

1.1 Model Description

In this work we use the MPC (Massively Parallel Computation) model, which can be traced back to descriptions given by Karloff et al. [19] and Feldman et al. [11], and was later refined by several works [1, 4, 15]. In the most general setting we have the following description: given a problem with input size N , there are M machines, each with S words of memory. The size of each machine’s memory is assumed to be considerably (e.g., polynomially) smaller than the entire data of the problem. More concretely, the memory size S is assumed to satisfy $S \leq N^{1-\alpha}$ for a constant $\alpha > 0$. Ideally, we want to be able to work with as small as possible memory requirement, per machine. Since the cluster of machines has to be able to store the input, a natural lower bound for the number of machines is $M \geq \frac{N}{S}$; this is typically tight up to a logarithmic factor and we usually assume $M = \tilde{\Theta}(\frac{N}{S})$. Initially the input is

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

SPAA '20, July 15–17, 2020, Virtual Event, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6935-0/20/07...\$15.00

<https://doi.org/10.1145/3350755.3400260>

divided arbitrarily among all machines. Computation proceeds in synchronous rounds, where in each round every machine can execute some computation on the data it holds. This local computation is restricted to be of polynomial running time with respect to the local memory size. After this computation, there is a round of communication where each machine can send to every other machine some data – thus the network graph is the complete graph. The only restriction on the communication is that the total amount of data that one machine sends or receives cannot exceed its memory capacity S . The bottleneck in this model is the communication. Therefore the analysis of an algorithm running in the MPC model is focused on the number of rounds.

When considering graph problems, for most problems, the technical difficulty (and the round complexity) of the problem increases as the memory per machine decreases. We will soon discuss instances of this. Considering this effect, in the literature, there is a further distinction of the regimes of the MPC model based on how much memory a machine has relative to the number of vertices, n , in the graph:

- (A) *Strongly super-linear memory regime*: $S \geq n^{1+\beta}$, for a constant $\beta \in (0, 1)$
- (B) *Near-linear memory regime*: $S \in \tilde{\Theta}(n)$
- (C) *Strongly sub-linear memory regime*: $S \leq n^{1-\beta}$, for a constant $\beta \in (0, 1)$

The goal in the area is to obtain fast algorithms for as small as possible memory requirements, per machine.

1.2 State-of-the-Art

First of all, by positive simulation results from Karloff et al. [19] and Goodrich et al. [15] it is known that any CREW PRAM algorithm using $O(n^{2-2\epsilon})$ total memory, $O(n^{2-2\epsilon})$ processors and $t = t(n)$ time can be run in $O(t)$ rounds of MPC. The primary goal in the MPC model is to find algorithms that run strictly (and ideally significantly) faster than their PRAM counterparts. For the problems in focus in this paper— $(2 + \epsilon)$ -approximate Vertex Cover, Maximal Matching, and $(1 + \epsilon)$ -approximate Maximum Matching— $O(\log n)$ time algorithms are known through this simulation and via classic results in the PRAM and LOCAL model [18, 24]. The primary goal for these problems in the MPC setting is to obtain algorithms with time complexity much faster than $O(\log n)$, ideally with as small as possible memory requirements per machine.

For the strongly super-linear memory regime, Lattanzi et al. [21] provided a constant round algorithm to find a maximal matching. Using that algorithm as a building block they also provided constant round algorithms to compute an 8-approximate weighted maximum matching, 2-approximate minimum vertex cover and 3/2-approximate minimum edge cover.

The near-linear memory regime presented much more difficulty and for a number of years, there was no sub-logarithmic time algorithm known for any of these problems. That changed with a breakthrough from Czumaj et al. that showed the first sub-logarithmic time algorithm for maximum matching [8]. In that work the authors present an algorithm that computes a $(2 + \epsilon)$ -approximate maximum matching in $O((\log \log n)^2)$ MPC rounds. Later this result has been improved and simplified by Assadi et al. [2] and Ghaffari et al. [13] to reach an $O(\log \log n)$ round algorithm that computes

a $(1 + \epsilon)$ -approximate maximum matching. Moreover, the method of Ghaffari et al. [13] also provides a $(2 + \epsilon)$ -approximate minimum cardinality vertex cover. Gamlath et al. [12] showed an extension of the matching problem to the weighted case, providing a $(1 + \epsilon)$ -approximate maximum weight matching in $O(\log \log n)$ rounds. However, the weighted case of the vertex cover problem has remained open and the best known algorithm remains at the $O(\log n)$ complexity that follows from the PRAM literature.

A remark regarding the strongly sub-linear memory regime.

The case of the strongly sub-linear memory regime appears to be considerably harder and there is no known $\text{poly}(\log \log n)$ round algorithm for the general case of any of the above problems. The best known result is a work of Ghaffari and Uitto [14] that provides $\tilde{O}(\sqrt{\log n})$ -round algorithms for maximal independent set, maximal matching, 2-approximation of minimum vertex cover, and $(1 + \epsilon)$ -approximation of maximum matching. For special graph families, concretely trees and low arboricity graphs, $\text{poly}(\log \log n)$ round algorithms were provided by Behnezhad et al. [5].

1.3 Our Contributions

Our main result, which positively answers the open question for the weighted case of vertex cover in the near-linear memory regime, is as follows:

THEOREM 1.1. *There is a randomized MPC algorithm, with $\tilde{O}(n)$ memory per machine, that computes a $(2 + \epsilon)$ -approximate minimum weight vertex cover in $O(\log \log d)$ rounds in any input graph with n vertices and average degree d , with high probability¹.*

We emphasize that our round complexity is a function of the average degree d , instead of the maximum degree Δ . We are not aware of any prior work in MPC where the round complexity is a function of the average degree.

Let us briefly discuss the method: Our algorithm follows a similar outline with the algorithm by Ghaffari et al. [13] for the unweighted case, and more generally the *round compression* idea introduced by Czumaj et al. [8], but with some crucial and important changes that allow us to handle the weighted case.

The general round compression technique works roughly as follows: randomly partition the vertices among a small set of machines and simulate many iterations of a suitable LOCAL algorithm on the induced subgraphs in these machines. Then communicate the results and repeat for a few steps. The power lies in the possibility to simulate up to a constant fraction of the iterations of the LOCAL algorithm, by just working on the randomly partitioned graph and without any further communication among the machines, since the error incurred from neglecting cross-partition interactions is very small (which can be shown using concentration inequalities).

Compared to the unweighted case, it is more difficult to analyze the behaviour of the progress that is being made during the algorithm, simply because there is an additional factor that influences the behaviour of a vertex, namely its weight. The behaviour of a vertex in the unweighted setting is completely dependent on its

¹As standard, we use the phrase *with high probability*, or the abbreviation w.h.p., to indicate that an event happens with probability at least $1 - \frac{1}{n^c}$ for any desirable constant $c > 0$.

degree, which is moreover possible to estimate after taking a random subgraph. This property does not hold when vertex weights are added—due to the inherent deviations in the related random variables—and hence additional ideas are needed to make sure that we can still simulate the vertices well and moreover, that our algorithm makes enough progress in each step.

After introducing the necessary background, in Section 3.2 we will give an overview of our novel ideas for resolving the above-mentioned issues in the weighted case. We suspect that some of the ideas presented there for handling weights in round compression might find applications in other instances of round compression, where we have to deal with weights (and where natural sampling ideas face deviation issues).

Implications for Congested Clique. A very closely related model, which has received significant attention over the past few years from the distributed computing community, is the *congested clique* model [23]. This model was initially proposed to capture computing on overlay networks. The network is presented as a fully connected graph, where on each node there is a machine and the machines can communicate in an all-to-all fashion. The precise method of communication is abstracted away – there could be a direct link between all machines or a routing protocol to make communication possible. Therefore, the communication is considered to be the bottleneck and only small messages can be sent (of size $O(\log n)$, with n the number of nodes in the graph). Computation proceeds again in synchronized rounds where each machine executes some local computation and then sends messages to other machines. The local memory and computation power are assumed to be unlimited.

In [6, Theorem 3.2] the authors show a two-way simulation to show that the near-linear memory MPC setting, in their words *semi-MapReduce*, is equivalent to congested clique. Thus, our result also implies an $O(\log \log d)$ round algorithm for $(2+\epsilon)$ approximate minimum weight vertex cover in the congested clique model.

Roadmap. The remaining sections are structured as follows: Section 2 presents some basic preliminaries and notation. In Section 3, we provide an in depth description of our algorithm. Then, Section 4 provides the detailed analysis of the memory requirements, round complexity, and accuracy achieved by our algorithm.

2 PRELIMINARIES

We denote an undirected graph $G = (V, E)$ by its vertex set V and its edge set E , where an edge $e \in E$ is an unordered pair of vertices, e.g. $e = (u, v)$ represents an undirected edge between the vertices $u, v \in V$. Furthermore, we denote the degree of a vertex with $d(v)$ and the maximum degree of any vertex present in the graph is Δ . In this paper we consider graphs with vertex weights, denoted $w(u) \in \mathbb{R}^+$ for $u \in V$. For a vertex subset $V' \subseteq V$, we will use the notation $E[V']$ to denote $\{e = (u, v) \in E \mid u \in V', v \in V'\}$, the edges in the induced subgraph of V' . We use $E[V'; V'']$ to denote $\{e = (u, v) \in E \mid u \in V', v \in V''\}$.

For asymptotic notations O, Θ, Ω , an additional tilde hides polylogarithmic factors. For example, $\tilde{O}(f)$ denotes $O(f \cdot \text{poly}(\log(f)))$.

We will frequently use the following form of Chernoff bounds to bound the tails of a sum of independent random variables.

$$\begin{array}{ll}
 \text{Primal} \\
 \min & \sum_{v \in V} z_v \cdot w(v) \\
 \text{s.t.} & z_u + z_v \geq 1 \quad \forall (u, v) \in E \\
 & z_v \geq 0 \quad \forall v \in V \\
 \\
 \text{Dual} \\
 \max & \sum_{e \in E} x_e \\
 \text{s.t.} & \sum_{e \ni v} x_e \leq w(v) \quad \forall v \in V \\
 & x_e \geq 0 \quad \forall e \in E
 \end{array}$$

Figure 1: Linear programming relaxation for MWVC

THEOREM 2.1 (CHERNOFF BOUNDS). *Let $X = \sum_{i=1}^n X_i$ be the summation of independent random variables, each assuming values in $[0, 1]$. Let $\mu = \mathbb{E}(X)$. Then*

- $\mathbb{P}(|X - \mu| \geq \delta\mu) \leq 2 \exp(-\delta^2\mu/3)$ for $0 \leq \delta \leq 1$.
- $\mathbb{P}(|X - \mu| \geq \delta\mu) \leq 2 \exp(-\delta\mu/3)$ for $\delta > 1$.

More information about the background and applications of these Chernoff bounds can be found in [10].

3 THE ALGORITHM

As discussed earlier in the introduction, our MPC algorithm follows the framework of [13] and uses the powerful *round compression* technique first introduced in [8]. Recall that a critical part of such an MPC algorithm is a centralized/LOCAL algorithm that allows for efficient simulation under random sampling.

In the following section we describe the centralized algorithm we will use and analyse its approximation guarantee. Then in the next section we give an overview of our MPC algorithm and highlight the differences with the algorithm for the unweighted case [13]. In the last section we outline the full MPC algorithm.

3.1 Centralized Algorithm

We first describe a centralized algorithm for computing a $(2 + \epsilon)$ -approximate weighted vertex cover, using the standard primal-dual framework; this approach can be traced back to the first studies on approximating the vertex cover problem [3, 16].

We maintain the dual variables $\{x_e\}_{e \in E}$ which form a valid *fractional matching*, i.e., they satisfy the dual constraints $\sum_{e \ni v} x_e \leq w(v)$ for all $v \in V$. Every vertex has a status of being *active* or *frozen*, indicating whether this vertex is still participating in the algorithm. We say an edge is *active*, if and only if both of its endpoints are active. We start with a valid fractional matching, and set all vertices to be *active*. Then, we slowly increase the dual variable x_e of every active edge e , while not violating the dual constraints. When the dual constraint of a vertex becomes near-tight, we freeze this vertex and include it in our vertex cover solution. In the end, we can show by weak LP-duality that this solution is indeed a $(2 + O(\epsilon))$ -approximation.

Algorithm 1 implements such a primal-dual scheme. The choices of the initial weights x_e and the thresholds $\mathcal{T}_{v,t}$ are not specified in the description. We will specify them later when we simulate this centralized algorithm in the MPC model and compare their behaviour. Next, we analyze the approximation guarantee of this centralized algorithm.

Algorithm 1: A generic centralized MWVC algorithm

- (1) Input: graph $G = (V, E)$, weight function $w : V \rightarrow \mathbb{R}^+$
- (2) Initialization: let $\{x_{e,0}\}$ ($x_{e,0} > 0$ for all $e \in E$) be an arbitrary valid fractional matching
- (3) Let $\mathcal{T}_{v,t}$ be arbitrary numbers from interval $[1 - 4\epsilon, 1 - 2\epsilon]$, for all $v \in V$ and integers $t \geq 0$
- (4) While at least one edge is active, iterate $t \leftarrow 0, 1, \dots$:
 - (a) For each active vertex v satisfying $y_{v,t} := \sum_{e \ni v} x_{e,t} \geq \mathcal{T}_{v,t} \cdot w(v)$: freeze v and its incident edges
 - (b) For each active edge e : $x_{e,t+1} := x_{e,t} / (1 - \epsilon)$
 - (c) For each frozen edge e : $x_{e,t+1} := x_{e,t}$
- (5) Return all frozen vertices as a vertex cover

Observation 3.1. For all $t \geq 0$, the dual constraint $\sum_{e \ni v} x_{e,t} \leq w(v)$ is satisfied for all $v \in V$. In other words, Algorithm 1 maintains a valid fractional matching.

PROOF. We do a proof by induction on t . The validity for $t = 0$ is ensured by the initialization requirement. For the inductive step, assume that $\{x_{e,t}\}_{e \in E}$ is a valid fractional matching. By the start of iteration $t + 1$, for any active vertex v ,

$$\sum_{e \ni v} x_{e,t+1} \leq \sum_{e \ni v} \frac{x_{e,t}}{1 - \epsilon} < \frac{\mathcal{T}_{v,t} w(v)}{1 - \epsilon} \leq \frac{(1 - 2\epsilon)w(v)}{(1 - \epsilon)} < w(v).$$

where the second inequality follows from the fact that vertex v was not frozen in iteration t . For any frozen vertex v ,

$$\sum_{e \ni v} x_{e,t+1} = \sum_{e \ni v} x_{e,t} \leq w(v). \quad \square$$

LEMMA 3.2 (WEAK LP-DUALITY). *Let OPT be the total weight of the minimum weight vertex cover C^* of graph $G = (V, E)$, and $\{x_e\}_{e \in E}$ be any fractional matching of G . Then $OPT \geq \sum_{e \in E} x_e$.*

PROOF. Observe that

$$OPT = \sum_{v \in C^*} w(v) \geq \sum_{v \in C^*} \sum_{u: (u,v) \in E} x_{(u,v)} \geq \sum_{e \in E} x_e. \quad \square$$

PROPOSITION 3.3. *When Algorithm 1 terminates, it returns a vertex cover C which satisfies*

$$w(C) \leq (2 + 10\epsilon)OPT,$$

where OPT is the weight of a minimum weight vertex cover.

PROOF. We first claim that the returned set of frozen vertices forms a valid vertex cover. This follows from the fact that the algorithm only terminates when all edges have been frozen, i.e. when they contain at least one vertex that is frozen. Therefore the set of frozen vertices covers all the edges.

Next we will relate the weight of the fractional matching to the size of the vertex cover and use LP-duality to prove the claimed approximation ratio. Denote the value of the final fractional matching by $W_M = \sum_{e \in E} x_e$. For every vertex v in the returned vertex cover C ,

$$y_v = \sum_{e \ni v} x_e \geq \mathcal{T}_{v,t^*} w(v) \geq (1 - 4\epsilon)w(v),$$

where t^* is the iteration in which v became frozen. As each edge can be covered at most twice, once per endpoint, we have

$$2W_M = 2 \sum_{e \in E} x_e \geq \sum_{v \in C} \sum_{e \ni v} x_e \geq (1 - 4\epsilon)w(C),$$

where $w(C)$ denotes $\sum_{v \in C} w(v)$. Then, by weak LP-duality,

$$w(C) \leq \frac{2}{1 - 4\epsilon} W_M \leq \frac{2}{1 - 4\epsilon} OPT < (2 + 10\epsilon)OPT. \quad \square$$

3.2 Overview

Before presenting our MPC algorithm, let us briefly review the algorithm for the unweighted case from Ghaffari et al. [13].

A recap on the approach of Ghaffari et al. [13]. Their algorithm proceeds in *phases*, where the machines only communicate after each phase. At the start of a phase, the vertices are partitioned uniformly at random among $m = \sqrt{\delta}$ machines, where δ is an upper bound on the current maximum degree (in the induced subgraph of nonfrozen vertices). Next, each machine gathers the induced subgraph on the vertices it received and simulates the LOCAL primal-dual algorithm (Algorithm 1, with $w(v) = 1, \forall v \in V$) on this subgraph, where the initialization of the fractional matching is taken as $x_{e,0} = 1/n$. This LOCAL algorithm is simulated by only inspecting the *local* neighborhood of each vertex, i.e. the neighbors that landed on the *same machine*. When checking the dual constraints (Line (4a) of Algorithm 1), the algorithm uses the (scaled) total weight of incident edges from local neighbors, which is an unbiased estimate of the total incident weight on the full graph. These estimates are sharply concentrated, and hence through the use of *random thresholds* $\mathcal{T}_{v,t}$, with good probability, the behaviour of this simulation is very close to the behaviour of the LOCAL algorithm on the full graph for all iterations.² Their algorithm proceeds by reducing the maximum degree until $\delta < \text{poly log } n$, at which point the algorithm terminates in one more step solving the remaining instance (with only $\tilde{O}(n)$ edges) on one machine.

Our algorithm for the weighted case uses the framework of Ghaffari et al. [13], but has a few key differences. In the following, we give a high-level description of our new techniques.

Non-uniform initialization of edge weights. First of all, we use a different initialization of the LOCAL algorithm. Instead of using the standard initialization $x_{(u,v)} = 1/n$, we use $x_{(u,v)} := \min \left\{ \frac{w(v)}{d(v)}, \frac{w(u)}{d(u)} \right\}$.³ This non-standard initialization will be crucial to the analysis of our MPC algorithm. We first give a succinct analysis of this initialization in the centralized setting.

PROPOSITION 3.4. *The initialization $x_{(u,v),0} := \min \left\{ \frac{w(v)}{d(v)}, \frac{w(u)}{d(u)} \right\}$ is valid. Moreover, Algorithm 1 terminates after $O(\log \Delta)$ iterations under this initialization.*

PROOF. For every vertex v , $\sum_{e \ni v} x_{e,0} \leq d(v) \cdot \frac{w(v)}{d(v)} = w(v)$.

For the running time, consider any edge $e = (u, v)$ and w.l.o.g. assume its initial weight is $x_{e,0} = \frac{w(u)}{d(u)}$. If e is active after $\log_{(1/(1-\epsilon))}(\Delta)$

²We refer readers to [13, Section 4.2] for more intuition on their random thresholding technique and a discussion of its necessity.

³The actual initialization used in our MPC simulation is slightly different regarding the definition of residual degrees $d(u)$, for technical reasons (see Remark 4.2).

iterations, we have that $x_e \geq w(u)$, which violates the dual constraint and cannot happen. Hence, the algorithm terminates after $\log_{(1/(1-\epsilon))}(\Delta) \in O(\log \Delta)$ iterations. $\square$

One can see that the standard initial assignment $x_{e,0} = 1/n$ also yields a correct LOCAL algorithm for the weighted vertex cover problem, assuming the weights of vertices are rescaled so that $w(v) \geq 1$. However, the running time of this LOCAL algorithm would depend on the size of the maximum vertex weight: $O(\log(Wn))$ with $W = \max_{v \in V} w(v)$, which is undesirable.

One might want to use $x_{(u,v)} := \min\{\frac{w(v)}{\Delta}, \frac{w(u)}{\Delta}\}$ instead of $x_{(u,v)} := \min\{\frac{w(v)}{d(v)}, \frac{w(u)}{d(u)}\}$ for initialization. The former has smaller weights, and hence causes the primal-dual algorithm to make progress slower, though this difference is not obvious in the LOCAL model—the former initialization achieves the same time bound as stated in Proposition 3.4. However, when performing MPC simulation of the LOCAL algorithm, we could only achieve $O(\log \log \Delta)$ round complexity when using the former way of initialization. Using the latter one, we can achieve round complexity $O(\log \log d)$ (where d is the average degree) as claimed in the main result.

Analysis of progress via orienting edges. An integral part of applying round compression is that the graph gets sparsified, and hence we need to quantify this progress. In Ghaffari et al.’s unweighted algorithm [13], each nonfrozen edge has the same weight $x_t = (1/n)/(1-\epsilon)^t$, so the number of nonfrozen neighbors of any vertex is upper bounded by $1/x_t$, which gives a natural characterization of the sparsity of the remaining graph. For the weighted case, that characterization of progress does not hold anymore, due to our non-uniform initialization.

Instead, we will use an *orientation* argument. We orient every edge (u, v) from u to v if $\frac{w(u)}{d(u)} < \frac{w(v)}{d(v)}$. Then, since every edge e outward from vertex u has initial weight equal to $\frac{w(u)}{d(u)}$, we can give a natural upper bound on the *out-degree* of u , and analyze the out-degree shrinking over time. Although we do not have control on the (undirected) degrees of vertices, our upper bounds on the out-degrees still allow us to bound the total number of remaining edges and measure the progress of the algorithm.

Not simulating low degree vertices. Recall that the LOCAL algorithm needs to be simulated after taking a random sample of the vertices, and the accuracy of the simulation relies on necessary concentration bounds of the incident edge weight for each vertex. In the weighted case, low degree vertices can cause difficulties for proving such concentration, since (1) they can have big initial weights (due to our non-standard initialization), which introduce large deviation; and, (2) the sampling rate is not enough for their neighborhood. (For the algorithm of the unweighted case this is not an issue, since a low degree vertex cannot become frozen during the early stages of the LOCAL algorithm)

To alleviate the issue we divide the vertices in two classes at the start of a phase: in V^{high} are all the nonfrozen vertices with a *high degree*, defined as vertices v such that $d(v) \geq d^{0.95}$, where $d(v)$ is the degree of vertex v with respect to nonfrozen neighbors only, and $d := \frac{1}{n} \sum_{v \in V \text{ nonfrozen}} d(v)$.⁴ The other nonfrozen vertices of

low degree are called *inactive*, and are gathered in $V^{inactive}$. Then only the vertices in V^{high} are partitioned and simulated in this phase. In the analysis we will see that despite only simulating a subset of the vertices the algorithm still makes enough progress in one phase on reducing the average degree to reach an overall running time of $O(\log \log d)$ MPC rounds.

Other changes in our analysis. Since we need to deal with weights and degrees in our sampling/simulation arguments, the required concentration bounds are more delicate than the previous work [13].

To simplify some parts of the analysis, we make another modification of the algorithm. Recall that in Ghaffari et al.’s [13] algorithm, the (scaled) total incident weight of local neighbors was used as an estimate of the actual total incident weight on the full graph. This estimate is unbiased, and could have error on either of the two directions. One of them is easy to deal with, while the other one requires a much more difficult analysis (see the discussion in [13, Section 4.4.4] on “late-bad vertices”). In our algorithm, we simply introduce a bias term to the estimator (Line 2(g)i of Algorithm 2), so that with high probability it only has one-sided error comparing to the actual total incident weight. This will make the analysis easier when we compare the behaviour of our MPC simulation with the centralized algorithm in Section 4.3.

3.3 MPC Simulation

Our MPC algorithm is given in Algorithm 2. As described in the previous section our MPC algorithm consists of several phases: each execution of the while-loop at Line (2) is a phase, and the final execution of the centralized algorithm at Line (3) is the last phase of the algorithm. Each phase consists of several iterations, which are similar as in the centralized algorithm.

Vertices may become frozen in a phase. Once frozen, they will remain frozen throughout the rest of the algorithm. An edge is called frozen if and only if at least one of its endpoints is frozen. If edge e becomes frozen in a phase, then by the end of this phase—in particular, at Line (2h) and Line (2j)—it will be assigned a nonnegative weight x_e^{MPC} , which will never be changed in the following phases. Since the weights of frozen edges are already finalized, we will use *residual weight* $w'(v) = w(v) - \sum_{e \ni v \text{ frozen}} x_e^{MPC}$ as the weight of vertex v when we start the new phase; this makes the analysis cleaner. After the algorithm terminates, every edge $e \in E$ will be frozen and have a finalized edge weight x_e^{MPC} .

Per phase, we partition the high-degree vertices V^{high} uniformly at random between $m = \sqrt{d}$ machines, and the respective induced subgraphs are gathered on each machine; the sampling probability is chosen such that the size of the induced subgraph for one machine does not exceed its memory constraint $\tilde{O}(n)$. Subsequently, the centralized algorithm is simulated in each of the m induced subgraphs locally. When comparing with the thresholds $\mathcal{T}_{v,t}$, we use $\tilde{y}_{v,t}^{MPC}$ as a local estimator of the total incident weight of v .

After all machines have finished their local simulation, we assign an edge weight x_e^{MPC} to every edge $e = (u, v) \in E[V^{high}]$ —especially those cross-partition edges which did not participate in the local simulation—based on the earliest iteration where either of u, v became frozen during their respective local simulation. For frozen edges $e \in E[V^{high}]$, x_e^{MPC} are their finalized edge weights. Edges in $E[V^{inactive}; V^{high}]$ may have got frozen, too; their weights

⁴Note that d is not quite the “average” degree; the denominator is always n regardless of the number of nonfrozen vertices v .

Algorithm 2: MPC-Simulation for MWVC

-
- (1) Input: graph $G = (V, E)$, weight function $w : V \rightarrow \mathbb{R}^+$
 - (2) While $d := \frac{1}{n} \sum_{v \in V \text{ nonfrozen}} d(v) > \log^{30} n$:
 - (a) Let $V^{high} \leftarrow \{v \text{ nonfrozen} \mid d(v) \geq d^{0.95}\}$,
 $V^{inactive} \leftarrow \{v \text{ nonfrozen} \mid d(v) < d^{0.95}\}$
 - (b) Compute residual weights for all $v \in V^{high}$:
 $w'(v) \leftarrow w(v) - \sum_{e \ni v, \text{ frozen}} x_e^{MPC}$
 - (c) Initial edge weights for all $e = (u, v) \in E[V^{high}]$:
 $x_{e,0}^{MPC} := \min \left\{ \frac{w'(u)}{d(u)}, \frac{w'(v)}{d(v)} \right\}$
 - (d) Let $\mathcal{T}_{v,t}$ be independent random numbers uniformly chosen from $[1 - 4\epsilon, 1 - 2\epsilon]$, for all $v \in V^{high}$ and $0 \leq t < I$
 - (e) Set number of machines $m := \sqrt{d}$, and number of iterations $I := \frac{\log m}{10 \log 15}$
 - (f) Partition V^{high} into m sets $V_1, \dots, V_m$ by assigning each vertex to a machine independently and uniformly at random
 - (g) For each $i \in \{1, \dots, m\}$ in parallel, iterate
 $t \leftarrow 0, 1, \dots, I - 1$:
 - (i) For each active $v \in V_i$ satisfying $\tilde{y}_{v,t}^{MPC} := 2m^{-0.2} \cdot 15^t + m \cdot \sum_{e \ni v, e \in E[V_i]} x_{e,t}^{MPC} \geq \mathcal{T}_{v,t} \cdot w'(v)$: freeze v and its incident edges
 - (ii) For each active edge $e \in E[V_i]$: $x_{e,t+1}^{MPC} := x_{e,t}^{MPC} / (1 - \epsilon)$
 - (iii) For each frozen edge $e \in E[V_i]$: $x_{e,t+1}^{MPC} := x_{e,t}^{MPC}$
 - (h) For each $e = (u, v) \in E[V^{high}]$: $x_{e,t+1}^{MPC} \leftarrow x_{e,t}^{MPC} / (1 - \epsilon)^{t'}$, where $0 \leq t' < I$ is the earliest iteration in which either one of u, v was frozen; or $t' = I$ if both remain active
 - (i) For each active $v \in V^{high}$ satisfying $y_v^{MPC} := \sum_{e \ni v, e \in E[V^{high}]} x_e^{MPC} \geq w'(v)$: freeze v and its incident edges
 - (j) For each $e \in E[V^{inactive}; V^{high}]$: $x_e^{MPC} \leftarrow 0$
 - (k) Update residual degree for all nonfrozen v : $d(v) \leftarrow$ number of nonfrozen neighbors of v
 - (3) Directly run the centralized algorithm in one machine on the subgraph induced by nonfrozen vertices, with residual weights $w'(v) \leftarrow w(v) - \sum_{e \ni v, \text{ frozen}} x_e$
 - (4) Return all frozen vertices as a vertex cover
-

are finalized as 0. We also freeze the vertices whose sum of incident x_e^{MPC} is too big, so as to prevent them from having negative residual weight in the next phase. Once the average degree is below $\log^{30} n$, there are at most $n \log^{30} n \in \tilde{O}(n)$ edges left.⁵ Then, we move all edges into one machine, which executes the last iterations of the centralized algorithm.

4 ANALYSIS

In this chapter we provide an analysis of Algorithm 2, the MPC simulation. The analysis is split into three major parts. In Section 4.1 we address the memory constraints for the machines. In Section 4.2 we derive the round complexity by analyzing the degree reduction

in each phase. Finally, we turn the attention to the approximation ratio in Section 4.3.

4.1 Memory Constraint

Recall that V^{high} is divided into subsets $V_1, \dots, V_m$, where each vertex $v \in V^{high}$ is independently randomly assigned to one of the subsets. To simulate one phase, the i -th machine ($1 \leq i \leq m$) needs to store the subgraph induced by V_i , together with edge weights $x_{e,0}^{MPC}$ and vertex weights $w'(v)$. We do not need to store the random thresholds $\mathcal{T}_{v,t}$, as they can be sampled on the fly. In the following lemma we show that the induced graph of V_i with high probability contains at most $O(n)$ edges, so the necessary information can fit into one machine.

LEMMA 4.1. *At Line (2f), with high probability $|E[V_i]| \in O(n)$ hold for all $1 \leq i \leq m$.*

PROOF. Recall that $m = \sqrt{d}$, and $V_1, \dots, V_m$ is a random partition of V^{high} . Fixing any $i \in \{1, \dots, m\}$, define independent random variables $s_v \in [0, 1]$ for all $v \in V^{high}$ as: if $v \in V_i$ then $s_v := d(v)/n$; otherwise $s_v := 0$. Then $\mathbb{E}[\sum_{v \in V^{high}} s_v] = \sum_{v \in V^{high}} \frac{d(v)/n}{m} \leq \frac{d}{m} = \sqrt{d}$. By Chernoff bound, we have

$$\mathbb{P}\left[\sum_{v \in V^{high}} s_v > 2\sqrt{d}\right] \leq \exp(-\sqrt{d}/3).$$

Recall that $d > \log^{30} n$. Hence, with high probability we have

$$\sum_{v \in V_i} d(v) = n \cdot \sum_{v \in V^{high}} s_v \leq 2n\sqrt{d}.$$

For any $v \in V^{high}$, let $d_i(v)$ denote the number of its neighbors in V_i . Similarly by Chernoff bound, we have

$$\mathbb{P}[d_i(v) > \frac{2d(v)}{m}] \leq \exp(-d(v)/3m) \leq \exp(-d^{0.45}/3),$$

where the last inequality follows from the definition of V^{high} . Hence, with high probability we have

$$|E(G'[V_i])| = \frac{1}{2} \sum_{v \in V_i} d_i(v) \leq \frac{1}{2} \sum_{v \in V_i} \frac{2d(v)}{m} \leq \frac{2n\sqrt{d}}{m} = 2n.$$

We finish the proof by a union bound over all $i \in \{1, \dots, m\}$. $\square$

We have shown that the near-linear local memory constraint is satisfied. Since the number of machines used for simulation is $m = \sqrt{d} \leq \sqrt{|E|/n}$, the total memory used is with high probability $\tilde{O}(\sqrt{dn}) \leq \tilde{O}(|E|)$, so the global memory constraint is also satisfied.

4.2 Round Complexity

Our algorithm runs in multiple phases, each of which can be implemented in $O(1)$ MPC rounds. The number of nonfrozen edges is reduced in each phase, and in the end we switch to the centralized algorithm when the nonfrozen edges fit in one machine. We will bound the round complexity by showing that the number of nonfrozen edges significantly decreases in each phase.

As discussed in Section 3.2, we use an orientation argument. At the beginning of the phase, we orient the edges $e \in E[V^{high}]$ in the following way: direct the edge (u, v) from u to v if $\frac{w'(u)}{d(u)} < \frac{w'(v)}{d(v)}$ and reverse otherwise, breaking ties arbitrarily. After this

⁵We did not attempt to optimize this 30 constant in the exponent.

orientation, the incident edges around vertex v split in two parts: $N_{in}(v)$ contains all edges directed towards v and $N_{out}(v)$ contains all edges directed outward from v . For each edge $e \in N_{out}(v)$ we have $x_{e,0}^{MPC} = \frac{w'(v)}{d(v)}$, and for each edge $e \in N_{in}(v)$ we have $x_{e,0}^{MPC} \leq \frac{w'(v)}{d(v)}$.

REMARK 4.2. Note that $d(v)$ is defined as the number of nonfrozen neighbors of v (see Line (2k)), i.e., the degree of v in the subgraph induced by $V^{high} \cup V^{inactive}$. It is not defined as the number of v 's neighbors in V^{high} .

A first observation to make is that the active out-degree decreases significantly over a phase.

Observation 4.3 (Active out-degree). After Line (2i) finishes, for any active vertex $v \in V^{high}$, denote with $d_A^{out}(v)$ the number of edges (v, u) directed outward from v such that $u \in V^{high}$ is still active. Then

$$d_A^{out}(v) \leq d(v)(1 - \varepsilon)^I.$$

PROOF. Recall that for every edge $e \in E[V^{high}]$ directed outward from v , we set $x_{e,0}^{MPC} = \frac{w'(v)}{d(v)}$ at the beginning of this phase. Assume towards a contradiction that after Line (2i) finishes there exists an active $v \in V^{high}$ with $d_A^{out}(v) > d(v)(1 - \varepsilon)^I$. The weight of active out-edges is at this point $x_e^{MPC} = x_{e,I}^{MPC} = \frac{w'(v)}{d(v)(1 - \varepsilon)^I}$. Therefore,

$$\begin{aligned} y_v^{MPC} &= \sum_{e \ni v; e \in E[V^{high}]} x_e^{MPC} \\ &\geq d_A^{out}(v) \cdot \frac{w'(v)}{d(v)(1 - \varepsilon)^I} \\ &> w'(v), \end{aligned}$$

meaning that v would have been frozen at Line (2i), hence there is no such vertex. $\square$

Using Observation 4.3 we show that the number of nonfrozen edges is decreasing over the course of a phase.

LEMMA 4.4. *With high probability, after Line (2k) finishes, the number of remaining nonfrozen edges is*

$$\frac{1}{2} \sum_{v \in V \text{ nonfrozen}} d(v) \leq 2nd(1 - \varepsilon)^I.$$

PROOF. A remaining nonfrozen edge is either incident to a low-degree vertex $v \in V^{inactive}$, or in $E[V^{high}]$ and has been active during this phase. Recall that

$$nd = \sum_{V^{high} \cup V^{inactive}} d(v)$$

(see Line (2) and Line (2a)). Combining Observation 4.3 and the definition of $V^{inactive}$, we have that the number of remaining active edges is at most

$$\begin{aligned} \sum_{v \in V^{high}} d_A^{out}(v) + n \cdot d^{0.95} &\leq \sum_{v \in V^{high}} d(v)(1 - \varepsilon)^I + nd^{0.95} \\ &\leq nd(1 - \varepsilon)^I + nd^{0.95} \\ &\leq 2nd(1 - \varepsilon)^I. \end{aligned}$$

The last inequality holds as $(1 - \varepsilon)^I = (1 - \varepsilon)^{\log d / 20 \log 15} \geq d^{-1/20}$. $\square$

THEOREM 4.5. *The MPC simulation with high probability takes at most $O(\log \log d)$ rounds, where $d = 2|E|/n$ is the initial average degree of the input graph.*

PROOF. As each phase takes $O(1)$ MPC rounds, we now need to show that the condition at Line (2) breaks after $O(\log \log d)$ phases. Using Lemma 4.4 and the definition of I (see Line (2e)), it remains to show that $d_k \leq \log^{30} n$ holds for some $k \in O(\log \log d)$, where

$$d_0 = d \text{ and } d_{i+1} \leq 4d_i(1 - \varepsilon)^{\frac{\log d_i}{20 \log 15}}.$$

Assume $0 < \varepsilon < 1/2$ and define the constant $\gamma = \frac{\log(1/(1-\varepsilon))}{40 \log 15} \in (0, 1)$. When $d_i > \log^{30} n$, we have

$$d_{i+1} \leq 4d_i^{1-2\gamma} \leq d_i^{1-\gamma}$$

for sufficiently large n . Hence we have $d_k \leq d_{k-1}^{1-\gamma} \leq \dots \leq d^{(1-\gamma)^k} \leq \log^{30} n$ for some $k \leq \frac{\log(\log d / 30 \log \log n)}{\log(1/(1-\gamma))} \in O(\log \log d)$. $\square$

4.3 Approximation Ratio

In this section, we will prove the approximation guarantee of our MPC algorithm. To achieve this, we consider one phase of the MPC algorithm, and imagine running the centralized algorithm on the induced subgraph of V^{high} , with the same vertex weights $w'(v)$, initial edge weights $x_{e,0}^{MPC}$ and random thresholds $\mathcal{T}_{v,t}$ that are used in this phase of MPC simulation, and then compare the behaviour of these two algorithms.

Recall that the edge weights and total incident weights in the centralized algorithm are denoted by $x_{e,t}$ and $y_{v,t}$. We have $x_{e,0} := x_{e,0}^{MPC}$. By Proposition 3.4 we know this initialization is valid.

In the description of the MPC algorithm, we have defined edge weights $x_{e,t}^{MPC}$ ($0 \leq t \leq I$) for all local edges $e \in E[V_1] \cup \dots \cup E[V_m]$. We can easily extend this definition to all edges in $E[V^{high}]$ (that is, including cross-partition edges), in the sense as given by Line (2h). Similarly we can define $y_{v,t}^{MPC}$ for all $v \in V^{high}$, $0 \leq t \leq I$ (as at Line (2i)).

We will prove the following key lemma, stating that the behaviour of our MPC simulation is similar to that of the centralized algorithm.

LEMMA 4.6. *Consider one phase of the MPC algorithm. Run the centralized algorithm for $I := \frac{\log m}{10 \log 15}$ iterations on the graph induced by V^{high} , with the same vertex weights w' , initial edge weights $x_{e,0}^{MPC}$, and random thresholds $\mathcal{T}_{v,t}$. With high probability, we have*

$$|y_{v,t} - \tilde{y}_{v,t}^{MPC}| \leq 6\varepsilon \cdot w'(v),$$

and

$$|y_{v,t} - y_{v,t}^{MPC}| \leq 6\varepsilon \cdot w'(v),$$

for all $0 \leq t \leq I$ and all $v \in V^{high}$.

Before proving this lemma, we show how it implies the approximation guarantee of our MPC algorithm.

THEOREM 4.7. *With high probability, Algorithm 2 returns a vertex cover C which satisfies $w(C) \leq (2 + 30\varepsilon)OPT$, where OPT is the weight of a minimum weight vertex cover.*

PROOF. The returned vertex cover consists of the vertices that were frozen in our MPC algorithm. When vertex v is frozen at Line (2(g)i) in some phase, we have $\tilde{y}_{v,t}^{MPC} \geq \mathcal{T}_{v,t} w'(v) \geq (1 - 4\epsilon)w'(v)$. By Lemma 4.6, this implies $y_{v,t}^{MPC} \geq (1 - 16\epsilon)w'(v)$. Recall that $w(v) - w'(v) \geq 0$ is the total weight of v 's incident edges that were frozen in previous phases. So we have

$$\sum_{e \ni v} x_e^{MPC} = w(v) - w'(v) + y_{v,t}^{MPC} \geq (1 - 16\epsilon)w(v).$$

Note that this inequality holds as well for those vertices v that became frozen at Line (2i) or in the final centralized phase (Line (3)).

On the other hand, by Observation 3.1, $y_{u,t} \leq w'(u)$ always holds for all vertices $u \in V^{high}$. Then by Lemma 4.6, we have $y_{u,t}^{MPC} \leq (1 + 6\epsilon)w'(u)$ with high probability. Similarly, this implies

$$\sum_{e \ni u} x_e^{MPC} \leq (1 + 6\epsilon)w(u)$$

in the end.

Hence, $\{x_e^{MPC}/(1 + 6\epsilon)\}_{e \in E}$ is a valid fractional matching, and by the same argument as in the proof of Proposition 3.3, we have

$$w(C) \leq 2 \sum_{e \ni v} \frac{x_e^{MPC}}{1 - 16\epsilon} \leq \frac{2(1 + 6\epsilon)OPT}{1 - 16\epsilon} \leq (2 + 30\epsilon)OPT. \quad \square$$

The rest of this section is devoted to proving Lemma 4.6.

In a phase, we say that vertex v becomes *bad* if it gets frozen in the centralized algorithm and not in MPC simulation (or the other way around). Once bad, the vertex remains bad throughout the whole phase. If a vertex is not bad, we say it is *good*.

In order to bound the weight of bad vertices, we will show that the estimated values $\tilde{y}_{v,t}^{MPC}$ and $y_{v,t}^{MPC}$ stay close to the actual values $y_{v,t}$ during one phase of MPC simulation. Once we establish that, we can use Lemma 4.8 below to bound the total weight of adjacent vertices which turn bad in a particular iteration.

LEMMA 4.8. *Suppose $|y_{v,t} - \tilde{y}_{v,t}^{MPC}| \leq \sigma w'(v)$ holds for all vertices v that are active in both the centralized algorithm and MPC simulation. Then, v becomes bad in iteration t with probability at most σ/ϵ and independently of other vertices.*

PROOF. For a vertex v to become bad, the estimate $\tilde{y}_{v,t}^{MPC}$ has to be on the other side of the threshold as $y_{v,t}$. Call the effective threshold $T_{v,t} = \mathcal{T}_{v,t} w'(v)$. Notice that when $|\tilde{y}_{v,t}^{MPC} - T_{v,t}| > \sigma w'(v)$ it is not possible for vertex v to become bad. Therefore, only if $T_{v,t}$ falls in the interval of size $2\sigma w'(v)$ around $\tilde{y}_{v,t}^{MPC}$, v might become bad. $\mathcal{T}_{v,t}$ is chosen uniformly at random from an interval with size 2ϵ , this is equivalent to pick $T_{v,t}$ uniformly at random from the scaled interval with size $2\epsilon w'(v)$. Therefore the possibility of v becoming bad in iteration t is upper-bounded by $2\sigma w'(v)/(2\epsilon w'(v)) = \sigma/\epsilon$. $\square$

In order to compare $y_{v,t}$ and $\tilde{y}_{v,t}^{MPC}$, we introduce an intermediate quantity defined as follows.

Definition 4.9. For $v \in V^{high}$, let

$$\tilde{y}_{v,t} := 2m^{-0.2} \cdot 15^t + m \cdot \sum_{e \ni v; e \in E[V_i]} x_{e,t},$$

where $v \in V_i$. Note that $\tilde{y}_{v,0} = \tilde{y}_{v,0}^{MPC}$, since $x_{e,0} = x_{e,0}^{MPC}$.

Instead of directly calculating a bound on $|y_{v,t} - y_{v,t}^{MPC}|$ and $|\tilde{y}_{v,t} - \tilde{y}_{v,t}^{MPC}|$, we use a slightly different notion introduced below.

Definition 4.10 (Weight difference). Let

$$\text{diff}(v, t) := \sum_{e \ni v; e \in E[V^{high}]} |x_{e,t} - x_{e,t}^{MPC}|.$$

And, let

$$\text{diff}^{local}(v, t) := m \cdot \sum_{e \ni v; e \in E[V_i]} |x_{e,t} - x_{e,t}^{MPC}|,$$

where $v \in V_i$.

Note that $|y_{v,t} - y_{v,t}^{MPC}| \leq \text{diff}(v, t)$, and $|\tilde{y}_{v,t} - \tilde{y}_{v,t}^{MPC}| \leq \text{diff}^{local}(v, t)$.

Now we prove the concentration lemma, which immediately implies a bound on $\tilde{y}_{v,t} - y_{v,t}$.

LEMMA 4.11 (CONCENTRATION). *Let $1 \leq r \leq m^{1.2}$, and let $U \subseteq V^{high}$ be a random subset where each vertex is included with probability $1/r$ independently. For any $0 \leq t \leq I$ and any vertex $v \in V^{high}$, with high probability,*

$$\left| y_{v,t} - r \sum_{u \in U; (v,u) \in E[V^{high}]} x_{(v,u),t} \right| < m^{-0.2} \cdot w'(v).$$

PROOF. Denote $w_{out}(v) = \frac{w'(v)}{d(v)} / (1 - \epsilon)^t$. For every incident edge $e = (u, v)$, we have

$$0 \leq x_{e,t} \leq x_{e,0} / (1 - \epsilon)^t = \min \left\{ \frac{w'(u)}{d(u)}, \frac{w'(v)}{d(v)} \right\} / (1 - \epsilon)^t \leq w_{out}(v).$$

Let

$$X_v := \sum_{u \in U; (u,v) \in E[V^{high}]} x_{(u,v),t} / w_{out}(v),$$

which is the sum of independent random variables in interval $[0, 1]$.

By the definition of set U it is clear that

$$\mathbb{E}[X_v] = \frac{1}{r} \sum_{e \ni v; e \in E[V^{high}]} x_{e,t} / w_{out}(v).$$

By Observation 3.1, $\sum_{e \ni v; e \in E[V^{high}]} x_{e,t} \leq w'(v) \leq d(v)w_{out}(v)$. We then obtain

$$d(v) \geq r \mathbb{E}[X_v]. \quad (1)$$

Let

$$\delta := \frac{m^{0.65} \sqrt{d(v)}}{r \mathbb{E}[X_v]}.$$

We use Chernoff bound and analyze two cases:

- If $\delta < 1$,

$$\begin{aligned} \mathbb{P}[|X_v - \mathbb{E}[X_v]| \geq \delta \mathbb{E}[X_v]] &\leq 2 \exp(-\delta^2 \mathbb{E}[X_v]/3) \\ &= 2 \exp\left(-\frac{m^{1.3} d(v)}{3r^2 \mathbb{E}[X_v]}\right) \\ &\leq 2 \exp(-m^{0.1}/3), \end{aligned}$$

where the last inequality follows from (1) and $r \leq m^{1.2}$.

- If $\delta \geq 1$,

$$\begin{aligned} \mathbb{P}[|X_v - \mathbb{E}[X_v]| \geq \delta \mathbb{E}[X_v]] &\leq 2 \exp(-\delta \mathbb{E}[X_v]/3) \\ &= 2 \exp\left(-\frac{m^{0.65} \sqrt{d(v)}}{3r}\right) \\ &\leq 2 \exp(-m^{0.4}/3), \end{aligned}$$

where the last inequality follows from $d(v) \geq d^{0.95} = m^{1.9}$ for all $v \in V^{high}$ and $r \leq m^{1.2}$.

As $m = \sqrt{d} > \log^{15} n$, we conclude that with high probability $|X_v - \mathbb{E}[X_v]| < \delta \mathbb{E}[X_v]$, or equivalently

$$\begin{aligned} |y_{v,t} - r \sum_{u \in U; (v,u) \in E[V^{high}]} x_{(v,u),t}| &< r w_{out}(v) \cdot \delta \mathbb{E}[X_v] \\ &= \frac{m^{0.65}}{\sqrt{d(v)}} \cdot w'(v)/(1-\varepsilon)^t \\ &\leq m^{-0.3} \cdot w'(v)/(1-\varepsilon)^t \\ &\leq m^{-0.2} \cdot w'(v), \end{aligned}$$

where the last inequality follows from $(1/(1-\varepsilon))^t \leq (1/(1-\varepsilon))^I \leq m^{0.1}$. $\square$

COROLLARY 4.12. *For any $0 \leq t \leq I$ and any vertex $v \in V^{high}$, with high probability,*

$$(2 \cdot 15^t - 1)m^{-0.2}w'(v) \leq \tilde{y}_{v,t} - y_{v,t} \leq (2 \cdot 15^t + 1)m^{-0.2}w'(v).$$

PROOF. The proof directly follows from the definition of V_i , Definition 4.9, and Lemma 4.11. $\square$

Now we will use induction to show that, $y_{v,t}, \tilde{y}_{v,t}, y_{v,t}^{MPC}$, and $\tilde{y}_{v,t}^{MPC}$ stay close to each other, for every good vertex v .

LEMMA 4.13. *For every $0 \leq t \leq I$, the following hold with high probability for every $v \in V^{high}$ that is good by the start of iteration t :*

- (1) $\text{diff}(v, t) \leq m^{-0.2} \cdot 15^t \cdot w'(v)$,
- (2) $\text{diff}^{local}(v, t) \leq m^{-0.2} \cdot 15^t \cdot w'(v)$, and,
- (3) $0 \leq \tilde{y}_{v,t}^{MPC} - y_{v,t} \leq 4 \cdot m^{-0.2} \cdot 15^t \cdot w'(v)$.

PROOF. Note that Item (3) directly follows from Item (2) and Corollary 4.12:

$$\begin{aligned} \tilde{y}_{v,t}^{MPC} - y_{v,t} &= (\tilde{y}_{v,t}^{MPC} - \tilde{y}_{v,t}) + (\tilde{y}_{v,t} - y_{v,t}) \\ &\leq \text{diff}^{local}(v, t) + m^{-0.2}(2 \cdot 15^t + 1)w'(v) \\ &\leq m^{-0.2} \cdot 4 \cdot 15^t \cdot w'(v), \end{aligned}$$

and,

$$\begin{aligned} \tilde{y}_{v,t}^{MPC} - y_{v,t} &= (\tilde{y}_{v,t}^{MPC} - \tilde{y}_{v,t}) + (\tilde{y}_{v,t} - y_{v,t}) \\ &\geq -\text{diff}^{local}(v, t) + m^{-0.2}(2 \cdot 15^t - 1)w'(v) \\ &\geq 0. \end{aligned}$$

In the following we will prove Item (1) and Item (2) by induction. Assuming the statements hold for some $t \geq 0$, we bound $\text{diff}(v, t+1)$ and $\text{diff}^{local}(v, t+1)$ by analyzing the evolution of weight differences in iteration t , for all v that remain good after iteration t finishes.

Old bad vertices: If $|x_{e,t} - x_{e,t}^{MPC}| > 0$ for some $e = (v, u)$, then we must have $x_{e,t} = x_{e,0}/(1-\varepsilon)^{t_1}$, $x_{e,t}^{MPC} = x_{e,0}/(1-\varepsilon)^{t_2}$ for some $t_1 \neq t_2$, that is, e was frozen at different iterations in the centralized and the MPC algorithms, or e was frozen in one algorithm but is still active in the other. In the former case, we have $|x_{e,t} - x_{e,t}^{MPC}| = |x_{e,t+1} - x_{e,t+1}^{MPC}|$. In the latter case, we assume w.l.o.g. $t_1 < t_2 = t$,

and then

$$\begin{aligned} \frac{|x_{e,t+1} - x_{e,t+1}^{MPC}|}{|x_{e,t} - x_{e,t}^{MPC}|} &= \frac{|x_{e,0}/(1-\varepsilon)^{t_1} - x_{e,0}/(1-\varepsilon)^{t_2+1}|}{|x_{e,0}/(1-\varepsilon)^{t_1} - x_{e,0}/(1-\varepsilon)^{t_2}|} \\ &= \frac{1 - (1-\varepsilon)^{1+t_2-t_1}}{1 - (1-\varepsilon)^{t_2-t_1}} \\ &< 3, \end{aligned}$$

for small enough ε .

New bad vertices: We now analyze the edges $e = (v, u)$ such that $|x_{e,t} - x_{e,t}^{MPC}| = 0$. If $|x_{e,t+1} - x_{e,t+1}^{MPC}| > 0$, then it must be that u turns bad in iteration t (since v remains good), and we have

$$|x_{e,t+1} - x_{e,t+1}^{MPC}| = x_{e,t}/(1-\varepsilon) - x_{e,t} \leq 2\varepsilon x_{e,t}.$$

By Item (3) and Lemma 4.8, each u that was good by the start of iteration t turns bad with probability at most

$$4m^{-0.2} \cdot 15^t \cdot w'(v)/\varepsilon$$

independently. Hence, by Lemma 4.11, the total contribution of such $|x_{e,t+1} - x_{e,t+1}^{MPC}|$ in $\text{diff}(v, t+1)$ is with high probability at most

$$\begin{aligned} \sum_{\substack{e=(v,u) \in E[V^{high}]; \\ u \text{ turns bad}}} 2\varepsilon x_{e,t} &\leq 2\varepsilon \cdot \frac{4m^{-0.2} \cdot 15^t w'(v)}{\varepsilon} (y_{v,t} + m^{-0.2} w'(v)) \\ &\leq 12m^{-0.2} \cdot 15^t \cdot w'(v), \end{aligned}$$

where we used $y_{v,t} \leq w'(v)$ and assumed $m^{-0.2} \leq 1/2$. Similarly, their contribution in $\text{diff}^{local}(v, t+1)$ is with high probability at most

$$\begin{aligned} m \cdot \sum_{\substack{e=(v,u) \in E[V_i]; \\ u \text{ turns bad}}} 2\varepsilon x_{e,t} &\leq m \cdot 2\varepsilon \frac{4m^{-0.2} 15^t w'(v)}{m\varepsilon} (y_{v,t} + m^{-0.2} w'(v)) \\ &\leq 12m^{-0.2} \cdot 15^t \cdot w'(v). \end{aligned}$$

Finally, combining the effect of old bad vertices and new bad vertices, we have

$$\begin{aligned} &\text{diff}(v, t+1) \\ &= \sum_{e \ni v; e \in E[V^{high}]} |x_{e,t+1} - x_{e,t+1}^{MPC}| \\ &= \sum_{\substack{e \ni v; e \in E[V^{high}]; \\ |x_{e,t} - x_{e,t}^{MPC}| > 0}} |x_{e,t+1} - x_{e,t+1}^{MPC}| + \sum_{\substack{e \ni v; e \in E[V^{high}]; \\ |x_{e,t} - x_{e,t}^{MPC}| = 0}} |x_{e,t+1} - x_{e,t+1}^{MPC}| \\ &\leq \sum_{\substack{e \ni v; e \in E[V^{high}]; \\ |x_{e,t} - x_{e,t}^{MPC}| > 0}} 3|x_{e,t} - x_{e,t}^{MPC}| + \sum_{\substack{e=(v,u) \in E[V^{high}]; \\ u \text{ turns bad}}} 2\varepsilon x_{e,t} \\ &\leq 3 \text{diff}(v, t) + 12m^{-0.2} \cdot 15^t \cdot w'(v) \\ &\leq m^{-0.2} \cdot 15^{t+1} \cdot w'(v). \end{aligned}$$

Similarly we can show $\text{diff}^{local}(v, t+1) \leq m^{-0.2} \cdot 15^{t+1} \cdot w'(v)$. $\square$

REMINDER OF LEMMA 4.6. *With high probability, we have*

$$|y_{v,t} - \tilde{y}_{v,t}^{MPC}| \leq 6\varepsilon \cdot w'(v),$$

and

$$|y_{v,t} - y_{v,t}^{MPC}| \leq 6\varepsilon \cdot w'(v),$$

for all $0 \leq t \leq I$ and all $v \in V^{high}$.

PROOF. If v is good by the start of iteration t , by Lemma 4.13,

$$0 \leq \tilde{y}_{v,t}^{MPC} - y_{v,t} \leq 4 \cdot 15^t m^{-0.2} w'(v) \leq 4m^{-0.1} w'(v) \leq \varepsilon w'(v),$$

where we used $t \leq I = \log m / (10 \log 15)$, and assumed $4m^{-0.1} \leq \varepsilon$. And similarly we have

$$|y_{v,t} - \tilde{y}_{v,t}^{MPC}| \leq \text{diff}(v, t) \leq 15^t m^{-0.2} w'(v) \leq \varepsilon w'(v).$$

Otherwise, suppose v turned bad in iteration $t^* < t$. Because $\tilde{y}_{v,t^*}^{MPC} - y_{v,t^*} \geq 0$, it must be that v became frozen in MPC simulation while remained active in the centralized algorithm in iteration t^* . So

$$\tilde{y}_{v,t}^{MPC} = \tilde{y}_{v,t^*}^{MPC} \geq \mathcal{T}_{v,t^*} w'(v) \geq (1 - 4\varepsilon) w'(v),$$

Since v was good by the start of iteration t^* , we have

$$y_{v,t^*} \geq \tilde{y}_{v,t^*}^{MPC} - \varepsilon w'(v) \geq (1 - 5\varepsilon) w'(v).$$

Then by monotonicity of $x_{e,t}$ and Observation 3.1, we have $y_{v,t^*} \leq y_{v,t} \leq w'(v)$, implying that

$$|y_{v,t^*} - y_{v,t}| \leq 5\varepsilon w'(v).$$

Hence, we have

$$\begin{aligned} |y_{v,t} - \tilde{y}_{v,t}^{MPC}| &= |y_{v,t} - \tilde{y}_{v,t^*}^{MPC}| \\ &\leq |y_{v,t} - y_{v,t^*}| + |y_{v,t^*} - \tilde{y}_{v,t^*}^{MPC}| \\ &\leq 5\varepsilon w'(v) + \varepsilon w'(v) \\ &= 6\varepsilon w'(v). \end{aligned}$$

Similarly,

$$|y_{v,t} - y_{v,t}^{MPC}| \leq 6\varepsilon w'(v). \quad \square$$

ACKNOWLEDGMENTS

We are grateful to the reviewers of SPAA 2020 for their helpful comments.

The first author's work in this project was supported by funding from the European Research Council (ERC), under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 853109).

REFERENCES

- [1] Alexandr Andoni, Aleksandar Nikolov, Krzysztof Onak, and Grigory Yaroslavtsev. 2014. Parallel Algorithms for Geometric Graph Problems. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*. 574–583. <https://doi.org/10.1145/2591796.2591805>
- [2] Sepehr Assadi, MohammadHossein Bateni, Aaron Bernstein, Vahab Mirrokni, and Cliff Stein. 2019. Coresets Meet EDCS: Algorithms for Matching and Vertex Cover on Massive Graphs. In *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1616–1635. <https://doi.org/10.1137/1.9781611975482.98>
- [3] Reuven Bar-Yehuda and Shimon Even. 1981. A Linear-time Approximation Algorithm for the Weighted Vertex Cover Problem. *Journal of Algorithms* 2, 2 (1981), 198–203. [https://doi.org/10.1016/0196-6774\(81\)90020-1](https://doi.org/10.1016/0196-6774(81)90020-1)
- [4] Paul Beame, Paraschos Koutris, and Dan Suciu. 2017. Communication Steps for Parallel Query Processing. *J. ACM* 64, 6 (2017), 1–58. <https://doi.org/10.1145/3125644>
- [5] Soheil Behnezhad, Sebastian Brandt, Mahsa Derakhshan, Manuela Fischer, MohammadTaghi Hajiaghayi, Richard M. Karp, and Jara Uitto. 2019. Massively Parallel Computation of Matching and MIS in Sparse Graphs. In *Proceedings of the 38th ACM Symposium on Principles of Distributed Computing (PODC)*. 481–490. <https://doi.org/10.1145/3293611.3331609>
- [6] Soheil Behnezhad, Mahsa Derakhshan, and MohammadTaghi Hajiaghayi. 2018. Semi-MapReduce Meets Congested Clique. *arXiv preprint* 1802.10297 (2018). arXiv:1802.10297
- [7] Soheil Behnezhad, MohammadTaghi Hajiaghayi, and David G. Harris. 2019. Exponentially Faster Massively Parallel Maximal Matching. In *Proceedings of the 60th IEEE Symposium on Foundations of Computer Science (FOCS)*. 1637–1649. <https://doi.org/10.1109/FOCS.2019.00096>
- [8] Artur Czumaj, Jakub Łącki, Aleksander Mądry, Slobodan Mitrović, Krzysztof Onak, and Piotr Sankowski. 2019. Round Compression for Parallel Matching Algorithms. *SIAM J. Comput.* (2019), STOC18–1–STOC18–44. <https://doi.org/10.1137/18M1197655>
- [9] Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: Simplified Data Processing on Large Clusters. *Commun. ACM* 51, 1 (2008), 107–113. <https://doi.org/10.1145/1327452.1327492>
- [10] Devdatt Dubhashi and Alessandro Panconesi. 2009. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press.
- [11] Jon Feldman, Shanmugavelayutham Muthukrishnan, Anastasios Sidiropoulos, Cliff Stein, and Zoya Svitkina. 2010. On distributing symmetric streaming computations. *ACM Transactions on Algorithms* 6, 4 (2010), 1–19. <https://doi.org/10.1145/1824777.1824786>
- [12] Buddhima Gamlath, Sagar Kale, Slobodan Mitrović, and Ola Svensson. 2019. Weighted Matchings via Unweighted Augmentations. In *Proceedings of the 38th ACM Symposium on Principles of Distributed Computing (PODC)*. 491–500. <https://doi.org/10.1145/3293611.3331603>
- [13] Mohsen Ghaffari, Themis Gouleakis, Christian Konrad, Slobodan Mitrović, and Ronitt Rubinfeld. 2018. Improved Massively Parallel Computation Algorithms for MIS, Matching, and Vertex Cover. In *Proceedings of the 37th ACM Symposium on Principles of Distributed Computing (PODC)*. 129–138. <https://doi.org/10.1145/3212734.3212743>
- [14] Mohsen Ghaffari and Jara Uitto. 2019. Sparsifying Distributed Algorithms with Ramifications in Massively Parallel Computation and Centralized Local Computation. In *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1636–1653. <https://doi.org/10.1137/1.9781611975482.99>
- [15] Michael T. Goodrich, Nodari Sitchinava, and Qin Zhang. 2011. Sorting, Searching, and Simulation in the MapReduce Framework. In *Proceedings of the 22nd International Symposium on Algorithms and Computation (ISAAC)*. 374–383. https://doi.org/10.1007/978-3-642-25591-5_39
- [16] Dorit S. Hochbaum. 1982. Approximation Algorithms for the Set Covering and Vertex Cover Problems. *SIAM J. Comput.* 11, 3 (1982), 555–556. <https://doi.org/10.1137/0211045>
- [17] Michael Isard, Mihai Badiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. 2007. Dryad: Distributed Data-Parallel Programs from Sequential Building Blocks. In *Proceedings of the 2nd ACM SIGOPS/EuroSys European Conference on Computer Systems*. 59–72. <https://doi.org/10.1145/1272996.1273005>
- [18] Amos Israeli and Alon Itai. 1986. A Fast and Simple Randomized Parallel Algorithm for Maximal Matching. *Inform. Process. Lett.* 22, 2 (1986), 77–80. [https://doi.org/10.1016/0020-0190\(86\)90144-4](https://doi.org/10.1016/0020-0190(86)90144-4)
- [19] Howard Karloff, Siddharth Suri, and Sergei Vassilvskii. 2010. A Model of Computation for MapReduce. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 938–948. <https://doi.org/10.1137/1.9781611973075.76>
- [20] Christos Koufogiannakis and Neal E. Young. 2009. 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)*. 171–179. <https://doi.org/10.1145/1582716.1582746>
- [21] Silvio Lattanzi, Benjamin Moseley, Siddharth Suri, and Sergei Vassilvskii. 2011. Filtering: A Method for Solving Graph Problems in MapReduce. In *Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 85–94. <https://doi.org/10.1145/1989493.1989505>
- [22] Nathan Linial. 1992. Locality in Distributed Graph Algorithms. *SIAM J. Comput.* 21, 1 (1992), 193–201. <https://doi.org/10.1137/0221015>
- [23] Zvi Lotker, Boaz Patt-Shamir, Elan Pavlov, and David Peleg. 2005. Minimum-Weight Spanning Tree Construction in $O(\log \log n)$ Communication Rounds. *SIAM J. Comput.* 35, 1 (2005), 120–131. <https://doi.org/10.1137/s0097539704441848>
- [24] Zvi Lotker, Boaz Patt-Shamir, and Seth Pettie. 2008. Improved Distributed Approximate Matching. In *Proceedings of the 20th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 129–136. <https://doi.org/10.1145/1378533.1378558>
- [25] Tom White. 2012. *Hadoop: The definitive guide*. “O’Reilly Media, Inc.”.
- [26] James C. Wyllie. 1979. *The Complexity of Parallel Computations*. Ph.D. Dissertation. Cornell University. <https://hdl.handle.net/1813/7502>
- [27] Matei Zaharia, Mosharaf Chowdhury, Michael J. Franklin, Scott Shenker, and Ion Stoica. 2010. Spark: Cluster Computing with Working Sets. In *Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing*.