# Distributed Triangle Detection via Expander Decomposition 

Yi-Jun Chang<br>University of Michigan<br>cyijun@umich.edu

Seth Pettie<br>University of Michigan<br>pettie@umich.edu

Hengjie Zhang<br>IIIS, Tsinghua University<br>zhanghj15@mails.tsinghua.edu.cn


#### Abstract

We present improved distributed algorithms for triangle detection and its variants in the CONGEST model. We show that Triangle Detection, Counting, and Enumeration can be solved in $\tilde{O}\left(n^{1 / 2}\right)$ rounds. In contrast, the previous state-of-the-art bounds for Triangle Detection and Enumeration were $\tilde{O}\left(n^{2 / 3}\right)$ and $\tilde{O}\left(n^{3 / 4}\right)$, respectively, due to Izumi and LeGall (PODC 2017).

The main technical novelty in this work is a distributed graph partitioning algorithm. We show that in $\tilde{O}\left(n^{1-\delta}\right)$ rounds we can partition the edge set of the network $G=(V, E)$ into three parts $E=E_{m} \cup E_{s} \cup E_{r}$ such that - Each connected component induced by $E_{m}$ has minimum degree $\Omega\left(n^{\delta}\right)$ and conductance $\Omega(1 / \operatorname{polylog}(n))$. As a consequence the mixing time of a random walk within the component is $O(\operatorname{polylog}(n))$. - The subgraph induced by $E_{s}$ has arboricity at most $n^{\delta}$. - $\left|E_{r}\right| \leq|E| / 6$.

All of our algorithms are based on the following generic framework, which we believe is of interest beyond this work. Roughly, we deal with the set $E_{s}$ by an algorithm that is efficient for low-arboricity graphs, and deal with the set $E_{r}$ using recursive calls. For each connected component induced by $E_{m}$, we are able to simulate CONGESTED-CLIQUE algorithms with small overhead by applying a routing algorithm due to Ghaffari, Kuhn, and Su (PODC 2017) for high conductance graphs.


## 1 Introduction

We consider Triangle Detection problems in distributed networks. In the LOCAL model [34], which has no limit on bandwidth, all variants of Triangle Detection can be solved in exactly one round of communication: every vertex $v$ simply announces its neighborhood $N(v)$ to all neighbors. However, in models that take bandwidth into account, e.g., CONGEST, Triangle Detection becomes significantly more complicated. Whereas many graph optimization problems studied in the CONGEST model are intrinsically "global" (i.e., require at least diameter time) [ $2,11,12,14,15,20,26]$, Triangle Detection is somewhat unusual in that it can, in principle, be solved using only locally available information.

The CONGEST Model. The underlying distributed network is represented as an undirected graph $G=(V, E)$, where each vertex corresponds to a computational device, and each edge corresponds to a bi-directional communication link. We assume each $v \in V$ initially knows some global parameters such as $n=|V|, \Delta=\max _{v \in V} \operatorname{deg}(v)$, and $D=\operatorname{diameter}(G)$. Each vertex $v$ has a distinct $\Theta(\log n)$-bit identifier $\operatorname{ID}(v)$. The computation proceeds according to synchronized rounds. In each round, each vertex $v$ can perform unlimited local computation, and may send a distinct $O(\log n)$-bit message to each of its neighbors.

Throughout the paper we only consider the randomized variant of CONGEST. Each vertex is allowed to generate unlimited local random bits, but there is no global randomness.

The Congested Clique Model. The CONGESTED-CLIQUE model is a variant of CONGEST that allows all-to-all communication. Each vertex initially knows its adjacent edges and the set of vertex IDs, which we can assume w.l.o.g. is $\{1, \ldots,|V|\}$. In each round, each vertex transmits $n-1$ $O(\log n)$-bit messages, one addressed to each vertex in the graph.

Intuitively, the CONGEST model captures two constraints in distributed computing: locality and bandwidth, whereas the CONGESTED-CLIQUE model only focuses on the bandwidth constraint. This difference makes the two models behave very differently. For instance, the minimum spanning tree (MST) problem can be solved in $O(1)$ rounds in the CONGESTED-CLIQUE [23], but its round complexity is $\tilde{\Theta}(D+\sqrt{n})$ in CONGEST [35,37].

One of the main reasons that some problems can be solved efficiently in CONGESTED-CLIQUE is due to the routing algorithm of Lenzen [30]. As long as each vertex $v$ is the source and the destination of at most $O(n)$ messages, we can deliver all messages in $O(1)$ rounds. Using this routing algorithm [30] as a communication primitive, many parallel algorithms can be transformed to efficient CONGESTED-CLIQUE algorithms [6]. For example, consider the distributed matrix multiplication problem, where the input matrices are distributed to the vertices such that the $i$ th vertex initially knows the $i$ th row. The problem can be solved in the CONGESTED-CLIQUE model in $\tilde{O}\left(n^{1 / 3}\right)$ rounds over semirings, or $\tilde{O}\left(n^{1-(2 / \omega)+o(1)}\right)=o\left(n^{0.158}\right)$ rounds over rings [6].

Distributed Routing in Almost Mixing Time. A uniform lazy random walk moves a token around an undirected graph by iteratively applying the following process for some number of steps: with probability $1 / 2$ the token stays at the current vertex and otherwise it moves to a uniformly random neighbor. In a connected graph $G=(V, E)$, the stationary distribution of a lazy random walk is $\pi(u)=\operatorname{deg}(u) /(2|E|)$. Informally, the mixing time $\tau_{\text {mix }}(G)$ of a connected graph $G$ is the minimum number of lazy random walk steps needed to get within a negligible distance of the stationary distribution. Formally:

Definition 1 (Mixing time [16]). Let $p_{t}^{s}(v)$ be the probability that after $t$ steps of a lazy random walk starting at $s$, the walk lands at $v$. The mixing time $\tau_{\text {mix }}(G)$ is the minimum $t$ such that for all $s \in V$ and $v \in V$, we have $\left|p_{t}^{s}(v)-\pi(v)\right| \leq \pi(v) /|V|$.

Ghaffari, Kuhn, and Su [16] proved that if each vertex $v$ is the source and the destination of at most $O(\operatorname{deg}(v))$ messages, then all messages can be routed to their destinations in $\tau_{\text {mix }}(G) \cdot 2^{O(\sqrt{\log n \log \log n)}}$ rounds. The $2^{O(\sqrt{\log n \log \log n})}$ factor has recently been improved [17] to $2^{O(\sqrt{\log n})}$. The implication of this result is that many problems that can be solved efficiently in the CONGESTED-CLIQUE can also be solved efficiently in CONGEST, but only if $\tau_{\text {mix }}(G)$ is small. In particular, MST can be solved in $\tau_{\text {mix }}(G) \cdot 2^{O(\sqrt{\log n})}$ rounds in CONGEST [17]. This shows that the $\tilde{\Omega}(\sqrt{n})$ lower bound $[35,37]$ can be bypassed in networks with small $\tau_{\text {mix }}(G)$.

At this point, a natural question to ask is whether or not this line of research $[16,17]$ can be extended to a broader class of graphs (that may have high $\tau_{\text {mix }}(G)$ ), or even general graphs. The main contribution of this paper is to show that this is in fact doable, and based on this approach we improve the state-of-the-art algorithms for triangle detection, counting, and enumeration.

Graph Partitioning. It is well known that any graph can be decomposed into connected components of conductance $\Omega(\epsilon / \log n)$ (and hence poly $\left(\epsilon^{-1}, \log n\right)$ mixing time) after removing at most an $\epsilon$-fraction of the edges [5,33,39,41]. Moshkovitz and Shapira [31] showed that this bound is essentially tight. In particular, removing any constant fraction $\epsilon$ of the edges, the remaining components have conductance at most $O\left((\log \log n)^{2} / \log n\right)$.

A slightly weaker version of this graph partition can be constructed in near-linear time (for fixed $\epsilon$ ) in the sequential computation model [39]. Their algorithm uses random walks to explore the graph locally to find a cut with edge sparsity $O(1 / \log n)$. If the output cut is $S$, then the time spent is $\tilde{O}(\operatorname{Vol}(S)) .{ }^{1}$ By iteratively finding a sparse cut and removing it from the graph, in $\tilde{O}(|E|)$ time a graph partition is obtained in which all components have $\Omega(1 / \operatorname{poly} \log (n))$ conductance.

This graph partition and the idea of local graph exploration have found many applications, such as solving linear systems [39], unique games [5,36, 41], analysis of personalized PageRank [3], minimum cut [24], and property testing [18, 28].

In this work, we show that a variant of this graph partition can be constructed efficiently in the CONGEST model. The new twist is to partition the edge set in three parts, rather than two (i.e., removed and remaining edges).

Distributed Triangle Detection. Many variants of the triangle detection problem have been studied in the literature [6,21].

Triangle Detection. Each vertex $v$ reports a bit $b_{v}$, and $\bigvee_{v} b_{v}=1$ if and only if the graph contains a triangle.

Triangle Counting. Each vertex $v$ reports a number $t_{v}$, and $\sum_{v} t_{v}$ is exactly the total number of triangles in the graph.

Triangle Enumeration. Each vertex $v$ reports a list $L_{v}$ of triangles, and $\bigcup_{v} L_{v}$ contains exactly those triangles in the graph.

Local Triangle Enumeration. It may be desirable that every triangle be reported by one of the three participating vertices. It is required that $L_{v}$ only contain triangles involving $v$.

[^0]Dolev, Lenzen, and Peled [9, Remark 1] showed that Triangle Enumeration can be solved deterministically in $O\left(n^{1 / 3} / \log n\right)$ time in the CONGESTED-CLIQUE. Censor-Hillel et al. [6] presented an algorithm for Triangle Detection and Dounting in CONGESTED-CLIQUE that takes $\tilde{O}\left(n^{1-(2 / \omega)+o(1)}\right)=o\left(n^{0.158}\right)$ time via a reduction to matrix multiplication. Izumi and LeGall [21] showed that in CONGEST, the Detection and Enumeration problems can be solved in $\tilde{O}\left(n^{2 / 3}\right)$ and $\tilde{O}\left(n^{3 / 4}\right)$ time, respectively. They also proved that in both CONGEST and CONGESTED-CLIQUE, the Enumeration problem requires $\Omega\left(n^{1 / 3} / \log n\right)$ time, improving an earlier $\Omega\left(n^{1 / 3} / \log ^{3} n\right)$ bound of Pandurangan et al. [32]. Izumi and LeGall [21] proved a large separation between the complexity of the Enumeration and Local Enumeration problems. If triangles must be reported by a participating vertex, $\Omega(n / \log n)$ time is necessary (and sufficient) in CONGEST/CONGESTED-CLIQUE. More generally, the lower bound on Local Enumeration is $\Omega(\Delta / \log n)$ when the maximum degree is $\Delta$.

In this paper, we show that Triangle Detection, Enumeration, and Counting can be solved in $\tilde{O}\left(n^{1 / 2}\right)$ time in CONGEST. This result is achieved by a combination of our new distributed graph partition algorithm, the multi-commodity routing of [ 16,17 ], and a randomized version of the CONGESTED-CLIQUE algorithm for Triangle Enumeration of $[6,9]$. We also show that when the input graph has high conductance/low mixing time, that Triangle Enumeration can be solved even faster, in $O\left(\tau_{\text {mix }}(G) n^{1 / 3+o(1)}\right)$ time.

### 1.1 Technical Overview

Consider a graph $G=(V, E)$. For a vertex subset $S$, we write $\operatorname{Vol}(S)$ to denote $\sum_{v \in S} \operatorname{deg}(v)$. Note that by default the degree is with respect to the original graph $G$. We write $\bar{S}=V \backslash S$, and let $\partial(S)=E(S, \bar{S})$ be the set of edges $e=\{u, v\}$ with $u \in S$ and $v \in \bar{S}$. The sparsity (or conductance) of a cut $(S, \bar{S})$ is defined as $\Phi(S)=|\partial(S)| / \min \{\operatorname{Vol}(S), \operatorname{Vol}(\bar{S})\}$. The conductance $\Phi_{G}$ of a graph $G$ is the minimum value of $\Phi(S)$ over all vertex subsets $S$.

We have the following relation [22] between the mixing time $\tau_{\text {mix }}(G)$ and conductance $\Phi_{G}$ :

$$
\Theta\left(\frac{1}{\Phi_{G}}\right) \leq \tau_{\text {mix }}(G) \leq \Theta\left(\frac{\log n}{\Phi_{G}^{2}}\right)
$$

In particular, if the inverse of the conductance is $n^{o(1)}$, then the mixing time is also $n^{o(1)}$.

Our Graph Partition. We introduce a new, efficiently computable graph decomposition that partitions the edge set into three parts.

Definition 2. An $n^{\delta}$-decomposition of a graph $G=(V, E)$ is a tripartition of the edge set $E=$ $E_{m} \cup E_{s} \cup E_{r}$ satisfying the following conditions.
(a) Each connected component induced by $E_{m}$ has $O$ (poly $\log n$ ) mixing time, and each vertex in the component has $\Omega\left(n^{\delta}\right)$ incident edges in $E_{m}$. That is, for each vertex $v \in V$, either $\operatorname{deg}_{E_{m}}(v)=0$ or $\operatorname{deg}_{E_{m}}(v)=\Omega\left(n^{\delta}\right)$.
(b) $E_{s}=\bigcup_{v \in V} E_{s, v}$, where $E_{s, v}$ is a subset of edges incident to $v$ and $\left|E_{s, v}\right| \leq n^{\delta}$. We view $E_{s, v}$ as oriented away from $v$. The overall orientation on $E_{s}$ is acyclic, which certifies that $E_{s}$ has arboricity ${ }^{2}$ at most $n^{\delta}$. Each vertex $v$ knows $E_{s, v}$.
(c) $\left|E_{r}\right| \leq|E| / 6$.

[^1]Throughout the paper we assume $\delta \in(0,1)$ is a constant. The main difference between our graph partition and the ones in other works [39] is that we allow a set $E_{s}$ that induces a low arboricity subgraph. The purpose of having the set $E_{s}$ is to allow us to design an efficient CONGEST algorithm to construct the partition. In the sequential computation model, a common approach to find a graph partition is to iteratively find a vertex set $S$ with small $\Phi(S)=O(1 / \operatorname{polylog}(n))$, and then include the boundary edges $\partial(S)$ in the set $E_{r}$ and remove them from the current graph. The number of iterations can be as high as $\tilde{\Theta}(n)$ since we could have $|S|=\tilde{O}(1)$.

To reduce the number of iterations to at most $O\left(n^{1-\delta}\right)$, before we start to find $S$, we do a preprocessing step that removes low degree vertices in such a way that each vertex has degree at least $\Omega\left(n^{\delta}\right)$ in the remaining graph. This guarantees that $|S|=\Omega\left(n^{\delta}\right)$, and so the number of iterations can be upper bounded by $O\left(n^{1-\delta}\right)$, since the total number of vertices is $n$.

### 1.2 Additional Related Works

Drucker et al. [10] showed an $\Omega\left(\frac{n}{e^{\sqrt{\log n}} \log n}\right)$ lower bound for triangle detection in the broadcast CONGESTED-CLIQUE model, where each vertex can only broadcast one message to all other vertices in each round. In the CONGEST model, lower bounds for finding a triangles and other motifs (subgraphs) has been studied in $[1,10,19,25]$. The problem of detecting a $k$-cycle has an $\tilde{\Omega}(\sqrt{n})$ lower bound, for any even number $k \geq 4[10,25]$. Detecting a $k$-clique requires $\tilde{\Omega}(\sqrt{n})$ rounds for every $4 \leq k \leq \sqrt{n}$, and $\tilde{\Omega}(\sqrt{n} / k)$ rounds for every $k \geq \sqrt{n}[7]$.

Any one-round algorithm for the triangle membership problem (each vertex decides whether it belongs to a triangle) requires messages of $\operatorname{size} \Omega(\Delta \log n)$ [1], which meets the trivial upper bound. The distributed triangle detection problem has also been studied in the property testing setting in the CONGEST model [13].

Das Sarma et al. [38] first studied the distributed sparsest cut problem. Specifically, given two parameters $b$ and $\phi$, if there exists a cut of balance at least $b$ and conductance at most $\phi$, their algorithm outputs a cut of conductance at most $\tilde{O}(\sqrt{\phi})$ in $\tilde{O}((n+(1 / \phi)) / b)$ rounds. This result was later improved by Kuhn and Molla [27] to $\tilde{O}(D+1 /(b \phi)) .{ }^{3}$ Their algorithms are built upon techniques in [8].

The local graph clustering algorithm of Spielman and Teng [39] has been improved, both in terms of running time and the quality of the cuts discovered; see, e.g., [3, 4, 29, 40].

### 1.3 Organization

In Section 2 we present a new distributed algorithm for partitioning a graph into expanding subgraphs and a low-arboricity subgraph. A key subroutine for finding a sparse cut is described in Section 3. Section 4 presents Triangle Enumeration algorithms for both expanding graphs and general graphs. We conclude in Section 5 with a conjecture on the complexity of distributed graph partitioning.

## 2 Algorithm for Graph Partitioning

We first introduce some notation. Let $\operatorname{deg}_{H}(v)$ be the degree of $v$ in the subgraph $H$, or in the graph induced by edge/vertex set $H$. Let $V\left(E^{*}\right)$ be the set of vertices induced by the edge set

[^2]$E^{*} \subseteq E$. The strong diameter of a subgraph $H$ of $G$ is defined as $\max _{u, v \in H} \operatorname{dist}_{H}(u, v)$ and the weak diameter of $H$ is $\max _{u, v \in H} \operatorname{dist}_{G}(u, v)$.

The goal of this section is to prove the following theorem.
Theorem 1. Given a graph $G=(V, E)$ with $n=|V|$, we can find, w.h.p., an $n^{\delta}$-decomposition in $\tilde{O}\left(n^{1-\delta}\right)$ rounds in the CONGEST model.

The algorithm for Theorem 1 is based on repeated application of a black box algorithm $\mathcal{A}^{*}$, which is given a subgraph $G^{\prime}=\left(V^{\prime}, E^{\prime}\right)$ of the original graph $G=(V, E)$, where $V^{\prime}=V\left(E^{\prime}\right)$, $n^{\prime}=\left|V^{\prime}\right|$, and $m^{\prime}=\left|E^{\prime}\right|$. In $\mathcal{A}^{*}$, vertices may halt the algorithm at different times.

Specification of the Black Box. The goal of $\mathcal{A}^{*}$ is, given $G^{\prime}=\left(V^{\prime}, E^{\prime}\right)$, to partition $E^{\prime}$ into $E^{\prime}=E_{m}^{\prime} \cup E_{s}^{\prime} \cup E_{r}^{\prime}$ satisfying some conditions. The edge set $E_{m}^{\prime}$ is partitioned into $E_{m}^{\prime}=\bigcup_{i=1}^{t} \mathcal{E}_{i}$. We write $\mathcal{V}_{i}=V\left(\mathcal{E}_{i}\right)$ and $\mathcal{G}_{i}=\left(\mathcal{V}_{i}, \mathcal{E}_{i}\right)$, and define $S=V^{\prime} \backslash\left(\bigcup_{i=1}^{t} \mathcal{V}_{i}\right)$.
(C1) The vertex sets $\mathcal{V}_{1}, \ldots, \mathcal{V}_{t}, S$ are disjoint and partition $V^{\prime}$.
(C2) The edge set $E_{s}^{\prime}$ can be decomposed as $E_{s}^{\prime}=\bigcup_{v \in S} E_{s, v}^{\prime}$, where $E_{s, v}^{\prime}$ is a subset of edges incident to $v$, viewed as oriented away from $v$. This orientation is acyclic. For each vertex $v$ such that $E_{s, v}^{\prime} \neq \emptyset$, we have $\left|E_{s, v}^{\prime}\right|+\operatorname{deg}_{E_{m}^{\prime}}(v) \leq n^{\delta}$. Each vertex $v$ knows the set $E_{s, v}^{\prime}$.
(C3) Consider a subgraph $\mathcal{G}_{i}=\left(\mathcal{V}_{i}, \mathcal{E}_{i}\right)$. Vertices in $\mathcal{V}_{i}$ halt after the same number of rounds, say $K$. Exactly one of the following subcases will be satisfied.
(C3-1) All vertices in $\mathcal{V}_{i}$ have degree $\Omega\left(n^{\delta}\right)$ in the subgraph $\mathcal{G}_{i}$, each connected component of $\mathcal{G}_{i}$ has $O($ poly $\log n)$ mixing time, and $K=O($ poly $\log n)$. Furthermore, every vertex in $\mathcal{V}_{i}$ knows that they are in this sub-case.
(C3-2) $\left|\mathcal{V}_{i}\right| \leq n^{\prime}-\tilde{\Omega}\left(K n^{\delta}\right)$, and every vertex in $\mathcal{V}_{i}$ knows they are in this subcase.
(C4) Each vertex $v \in S$ halts in $\tilde{O}\left(n^{\prime} / n^{\delta}\right)$ rounds.
(C5) The inequality $E_{r}^{\prime} \leq\left(\left|E^{\prime}\right| \log \left|E^{\prime}\right|-\sum_{i=1}^{t}\left|\mathcal{E}_{i}\right| \log \left|\mathcal{E}_{i}\right|\right) /(6 \log m)$ is met.
(C6) Each cluster $\mathcal{V}_{i}$ has a distinct identifier. When a vertex $v \in \mathcal{V}_{i}$ terminates, $v$ knows the identifier of $\mathcal{V}_{i}$. If $v \in S, v$ knows that it belongs to $S$.

We briefly explain the intuition behind these conditions. The algorithm $\mathcal{A}^{*}$ will be applied recursively to all subgraphs $\mathcal{G}_{i}$ that have yet to satisfy the minimum degree and mixing time requirements specified in Theorem 1 and Definition 2. Because vertices in different components halt at various times, they also may begin these recursive calls at different times.

The goal of (C2) is to make sure that once a vertex $v$ has $E_{s, v}^{\prime} \neq \emptyset$, the total number of edges added to $E_{s, v}$ cannot exceed $n^{\delta}$. The goal of (C3) is to guarantee that the component size drops at a fast rate. The idea of (C5) is that the size of $E_{r}^{\prime}$ can be mostly charged to the number of the edges in the small-sized edge sets $\mathcal{E}_{i}$; this is used to bound the size of $E_{r}$ of our graph partitioning algorithm.

Note that in general the strong diameter of a subgraph $\mathcal{G}_{i}$ can be much higher than the maximum running time of vertices in $\mathcal{G}_{i}$, and it could be possible that $\mathcal{G}_{i}$ is not even a connected subgraph of $G$. However, (C6) guarantees that each vertex $v \in \mathcal{V}_{i}$ still knows that it belongs to $\mathcal{V}_{i}$. This property allows us to recursively execute $\mathcal{A}^{*}$ on each subgraph $\mathcal{G}_{i}$.

Lemma 2.1. There is an algorithm $\mathcal{A}^{*}$ that finds a partition $E^{\prime}=E_{m}^{\prime} \cup E_{s}^{\prime} \cup E_{r}^{\prime}$ meeting the above specification in the CONGEST model, w.h.p.

Assuming Lemma 2.1, we are now in a position to prove Theorem 1.
Proof of Theorem 1. Let $\mathcal{A}^{*}$ be the algorithm for Lemma 2.1. Initially, we apply $\mathcal{A}^{*}$ with $G^{\prime}=G$, and this returns a partition $E^{\prime}=E_{m}^{\prime} \cup E_{s}^{\prime} \cup E_{r}^{\prime}$.

For each subgraph $\mathcal{G}_{i}$ in the partition output by an invocation of $\mathcal{A}^{*}$, do the following. If $\mathcal{G}_{i}$ satisfies (C3-1), by definition it must have $O$ (poly $\log n$ ) mixing time, and all vertices in $\mathcal{G}_{i}$ have degree $\Omega\left(n^{\delta}\right)$ in $\mathcal{G}_{i}$; we add the edge set $\mathcal{E}_{i}$ to the set $E_{m}$ and all vertices in $\mathcal{V}_{i}$ halt. Otherwise we apply the algorithm recursively to $\mathcal{G}_{i}$, i.e., we begin by applying $\mathcal{A}^{*}$ to $G^{\prime}=\mathcal{G}_{i}$ to further partition its edges. All recursive calls proceed in parallel, but may begin and end at different times. Conditions (C1) and (C6) guarantee that this is possible. (Note that if $\mathcal{G}_{i}$ is disconnected, then each connected component of $\mathcal{G}_{i}$ will execute the algorithm in isolation.)

Initially $E_{r}=\emptyset$ and $E_{s}=\emptyset$. After each invocation of $\mathcal{A}^{*}$, we update $E_{r} \leftarrow E_{r} \cup E_{r}^{\prime}, E_{s} \leftarrow E_{s} \cup E_{s}^{\prime}$, and $E_{s, u} \leftarrow E_{s, u} \cup E_{s, u}^{\prime}$ for each vertex $u$.

Analysis. We verify that the three conditions of Definition 2 are satisfied. First of all, note that each connected component of $E_{m}$ terminated in (C3-1) must have $O$ (poly $\log n$ ) mixing time, and all vertices in the component have degree $\Omega\left(n^{\delta}\right)$ within the component. Condition (a) of Definition 2 is met. Next, observe that Condition (b) of Definition 2 is met due to (C2). If the output of $\mathcal{A}^{*}$ satisfies that $E_{s, v}^{\prime} \neq \emptyset$, then $\left|E_{s, v}\right|$ together with the number of remaining incident edges (i.e., the ones in $\left.E_{m}^{\prime}\right)$ is less then $n^{\delta}$. Therefore, $\left|E_{s, v}\right|$ cannot exceed $n^{\delta}$, since only the edges in $E_{m}^{\prime}$ that are incident to $v$ can be added to $E_{s, v}$ in future recursive calls. Lastly, we argue that (C5) implies that Condition (c) of Definition 2 is satisfied. Assume, inductively, that a recursive call on edge set $\mathcal{E}_{i}$ eventually contributes at most $\left|\mathcal{E}_{i}\right| \log \left|\mathcal{E}_{i}\right| /(6 \log m)$ edges to $E_{r}$. It follows from (C5) that the recursive call on edge set $E^{\prime}$ contributes $\left|E^{\prime}\right| \log \left|E^{\prime}\right| /(6 \log m)$ edges to $E_{r}$. We conclude that $\left|E_{r}\right| \leq|E| \log |E| /(6 \log |E|)=|E| / 6$.

Now we analyze the round complexity. In one recursive call of $\mathcal{A}^{*}$, consider a component $\mathcal{G}_{i}$ in the output partition, and let $K$ be the running time of vertices in $\mathcal{V}_{i}$. Due to (C3), there are two cases. If $\mathcal{G}_{i}$ satisfied (C3-1), it will halt in $K=O$ (poly $\log n$ ) rounds. Otherwise, (C3-2) is met, and we have $\left|\mathcal{V}_{i}\right| \leq n^{\prime}-\tilde{\Omega}\left(K n^{\delta}\right)$. Let $v \in V$ be any vertex, and let $K_{1}, \ldots, K_{z}$ be the running times of all calls to $\mathcal{A}^{*}$ that involve $v$. (Whenever $v$ ends up in $S$ or in a component satisfying (C3-1) it halts permanently, so $K_{1}, \ldots, K_{z-1}$ reflect executions that satisfy (C3-2) upon termination.) Then we must have $\sum_{i=1}^{z} K_{i} \leq \tilde{O}\left(n / n^{\delta}\right)+O(\operatorname{poly} \log n)=\tilde{O}\left(n^{1-\delta}\right)$. Thus, the whole algorithm stops within $\tilde{O}\left(n^{1-\delta}\right)$ rounds.

### 2.1 Subroutines

Before proving Lemma 2.1, we first introduce some helpful subroutines. Lemma 2.3 shows that for subgraphs of sufficiently high strong diameter, we can find a sparse cut of the subgraph, with runtime proportional to the strong diameter. Lemma 2.4 offers a procedure that removes a set of edges in such a way that the vertices in the remaining graph have high degree, and the removed edges form a low arboricity subgraph. Lemma 2.5 shows that if a subgraph already has a low conductance cut, then we can efficiently find a cut of similar quality.

All these subroutines are applied to a connected subgraph $G^{*}=\left(V^{*}, E^{*}\right)$ of the underlying network $G=(V, E)$, and the computation does not involve vertices outside of $G^{*}$. In subsequent discussion in this section, the parameters $n$ and $m$ are always defined as $n=|V|$ and $m=|E|$, which are independent of the chosen subgraph $G^{*}$.

Lemma 2.2. Let $m$ and $D$ be two numbers. Let $\left(a_{1}, \ldots, a_{D}\right)$ be a sequence of positive integers such that $D \geq 48 \log ^{2} m$ and $\sum_{i=1}^{D} a_{i} \leq m$. Then there exists an index $j$ such that $j \in[D / 4,3 D / 4]$ and

$$
a_{j} \leq \frac{1}{12 \log m} \cdot \min \left(\sum_{i=1}^{j-1} a_{i}, \sum_{i=j+1}^{D} a_{i}\right) .
$$

Proof. Define $S_{k}=\sum_{i=1}^{k} a_{i}$ to be the $k$ th prefix sum. By symmetry, we may assume $S_{\lfloor D / 2\rfloor} \leq$ $S_{D}-S_{\lfloor D / 2\rfloor}$, since otherwise we can reverse the sequence. Scan each index $j$ from $D / 4$ to $D / 2$. If an index $j$ does not satisfy $a_{j} \leq \frac{1}{12 \log m} \cdot S_{j-1}$, then this implies that $S_{j}>S_{j-1}\left(1+\frac{1}{12 \log m}\right)$. If no index $j \in[D / 4, D / 2]$ satisfies this condition then $S_{\lfloor D / 2\rfloor}$ is larger than

$$
S_{\lfloor D / 4\rfloor} \cdot\left(1+\frac{1}{12 \log m}\right)^{D / 4} \geq S_{\lfloor D / 4\rfloor} \cdot\left(1+\frac{1}{12 \log m}\right)^{12 \log ^{2} m} \geq S_{\lfloor D / 4\rfloor} \cdot m
$$

which is impossible since $\sum_{i=1}^{D} a_{i} \leq m$. Therefore, there must exist an index $j \in[D / 4, D / 2]$ such that $a_{j} \leq \frac{1}{12 \log m} \cdot S_{j-1}=\frac{1}{12 \log m} \cdot \sum_{i=1}^{j-1} a_{i}$. By our assumption that $S_{\lfloor D / 2\rfloor} \leq S_{D}-S_{\lfloor D / 2\rfloor}$, we also have $a_{j} \leq \frac{1}{12 \log m} \cdot \min \left(\sum_{i=1}^{j-1} a_{i}, \sum_{i=j+1}^{D} a_{i}\right)$.

Lemma 2.3 (High Diameter subroutine). Let $G^{*}=\left(V^{*}, E^{*}\right)$ be a connected subgraph and $x \in V^{*}$ be a vertex for which $\tilde{D}=\max _{v \in V^{*}} \operatorname{dist}_{G^{*}}(x, v) \geq 48 \log ^{2} m$. Define $V_{\text {low }}=\left\{v \in V^{*} \mid \operatorname{deg}_{G^{*}}(v) \leq\right.$ $\left.n^{\delta} / 2\right\}$. Suppose there are no edges connecting two vertices in $V_{\text {low }}$. Then we can find a cut ( $\left.C, C\right)$ of $G^{*}$ such that $\min (|C|,|\bar{C}|) \geq \frac{\tilde{D}}{32} n^{\delta}$ and $\partial(C) \leq \min (\operatorname{Vol}(C), \operatorname{Vol}(\bar{C})) /(12 \log m)$ in $O(\tilde{D})$ rounds deterministically in the CONGEST model. Each vertex in $V^{*}$ knows whether or not it is in $C$.

Proof. The algorithm is as follows. First, build a BFS tree of $G^{*}$ rooted at $x \in V^{*}$ in $O(\tilde{D})$ rounds. Let $L_{i}$ be the set of vertices of level $i$ in the BFS tree, and let $p_{i}$ be the number of edges $e=\{u, v\}$ such that $u \in L_{i}$ and $v \in L_{i+1}$. We write $L_{a . . b}=\bigcup_{i=a}^{b} L_{i}$. In $O(\tilde{D})$ rounds we can let the root $x$ learn the sequence $\left(p_{1}, \ldots, p_{\tilde{D}}\right)$.

Note that in a BFS tree, edges do not connect two vertices in non-adjacent levels. By Lemma 2.2, there exists an index $j \in[\tilde{D} / 4,3 \tilde{D} / 4]$ such that $p_{j} \leq \frac{1}{12 \log m} \cdot \min \left(\operatorname{Vol}\left(L_{1 . . j}\right), \operatorname{Vol}\left(L_{j+1 . . \tilde{D}}\right)\right)$, and such an index $j$ can be computed locally at the vertex $x$.

The cut is chosen to be $C=L_{1 . . j}$, so we have $\partial(C) \leq \min (\operatorname{Vol}(C), \operatorname{Vol}(\bar{C})) /(12 \log m)$. As for the second condition, due to our assumption in the statement of the lemma, for any two adjacent levels $L_{i}, L_{i+1}$, there must exist a vertex $v \in L_{i} \cup L_{i+1}$ such that $v \notin V_{\text {low }}$. By definition of $V_{\text {low }}, v$ has more than $n^{\delta} / 2$ neighbors in $G^{*}$, and they are all within $L_{i-1 . . i+2}$. Thus, the number of vertices within any four consecutive levels must be greater than $n^{\delta} / 2$. Since $j \in[\tilde{D} / 4,3 \tilde{D} / 4]$, we have

$$
\min (|C|,|\bar{C}|) \geq \frac{\tilde{D}}{4} / 4 \cdot n^{\delta} / 2 \geq \frac{\tilde{D}}{32} n^{\delta} .
$$

To let each vertex in $V^{*}$ learn whether or not it is in $C$, the root $x$ broadcasts the index $j$ to all vertices in $G^{*}$. After that, each vertex in level smaller than or equal to $j$ knows that it is in $C$; otherwise it is in $\bar{C}$.

Intuitively, Lemma 2.4 says that after the removal of a subgraph of small arboricity (i.e., the edge set $E_{s}^{\diamond}$ ), the remaining graph (i.e., the edge set $E^{\diamond}$ ) has high minimum degree. The runtime is proportional to the number of removed vertices (i.e., $\left|V^{*}\right|-\left|V^{\diamond}\right|$ ) divided by the threshold $n^{\delta}$. Note that the second condition of Lemma 2.4 implies that $E_{s, v}^{\diamond}=\emptyset$ for all $v \in V^{\diamond}$.

Lemma 2.4 (Low Degree subroutine). Let $G^{*}=\left(V^{*}, E^{*}\right)$ be a connected subgraph with strong diameter $D$. We can partition $E^{*}=E^{\diamond} \cup E_{s}^{\diamond}$ meeting the following two conditions.

1. Let $V^{\diamond}$ be the set of vertices induced by $E^{\diamond}$. Each $v \in V^{\diamond}$ has more than $n^{\delta} / 2$ incident edges in $E^{\diamond}$.
2. The edge set $E_{s}^{\diamond}$ is further partitioned as $E_{s}^{\diamond}=\bigcup_{v \in V^{*} \backslash V^{\diamond}} E_{s, v}^{\diamond}$, where $E_{s, v}^{\diamond}$ is a subset of incident edges of $v$, and $\left|E_{s, v}^{\diamond}\right| \leq n^{\delta}$. Each vertex $v$ knows $E_{s, v}^{\diamond}$.

This partition can be found in $O\left(D+\left(\left|V^{*}\right|-\left|V^{\diamond}\right|\right) / n^{\delta}\right)$ rounds deterministically in the CONGEST model.

Proof. To meet Condition 1, a naive approach is to iteratively "peel off" vertices that have degree at most $n^{\delta} / 2$, i.e., put all their incident edges in $E_{s}$, so long as any such vertex exists. On some graphs this process requires $\Omega(n)$ peeling iterations.

We solve this issue by doing a batch deletion. First, build a BFS tree of $G^{*}$ rooted at an arbitrary vertex $x \in V^{*}$. We use this BFS tree to let $x$ count the number of vertices that have degree less than $n^{\delta}$ in the remaining subgraph in $O(D)$ rounds.

The algorithm proceeds in iterations. Initially we set $E^{\diamond} \leftarrow E^{*}$ and $E_{s}^{\diamond} \leftarrow \emptyset$. In each iteration, we identify the subset $Z \subseteq V^{*}$ whose vertices have at most $n^{\delta}$ incident edges in $E^{\diamond}$. We orient all the $E^{\diamond}$-edges touching $Z$ away from $Z$, if one endpoint is in $Z$, or away from the endpoint with smaller ID, if both endpoints are in $Z$. Edges incident to $v$ oriented away from $v$ are added to $E_{s, v}^{\diamond}$ and removed from $E^{\diamond}$. The root $x$ then counts the number $z=|Z|$ of such vertices via the BFS tree. If $z>n^{\delta} / 2$, we proceed to the next iteration; otherwise we terminate the algorithm.

The termination condition ensures that each vertex has degree at least $\left(n^{\delta}+1\right)-z>n^{\delta} / 2$, and so Condition 1 is met. It is straightforward to see that the set $E_{s}^{\diamond}$ generated by the algorithm meets Condition 2 , since for each $v$, we only add edges to $E_{s, v}^{\diamond}$ once, and it is guaranteed that $\left|E_{s, v}^{\diamond}\right| \leq n^{\delta}$. Tie-breaking according to vertex-ID ensures the orientation is acyclic.

Throughout the process, each time one vertex puts any edges into $E_{s}^{\diamond}$, it no longer stays in $V^{\diamond}$. Each iteration can be done in $O(D)$ time. We proceed to the next iteration only if there are more than $n^{\delta} / 2$ vertices being removed from $V^{\diamond}$. A trivial implementation can lead to an algorithm taking $\left.O\left(D\left[\left(\left|V^{*}\right|-\left|V^{\diamond}\right|\right) / n^{\delta}\right\rceil\right)\right)$ rounds. The round complexity can be further improved to $O\left(D+\left(\left|V^{*}\right|-\left|V^{\diamond}\right|\right) / n^{\delta}\right)$ by pipelining the iterations. At some point the root $x$ detects that iteration $i$ was the last iteration; in $O(D)$ time it broadcasts a message to all nodes instructing them to roll back iterations $i+1, i+2, \ldots$, which have been executed speculatively.

The proof of the following lemma is deferred to Section 3. It is a consequence of combining Lemmas 3.4 and 3.5.

Lemma 2.5 (Low Conductance subroutine). Let $G^{*}=\left(V^{*}, E^{*}\right)$ be a connected subgraph with strong diameter $D$. Let $\phi \leq 1 / 12$ be a number. Suppose that there exists a subset $S \subset V^{*}$ satisfying

$$
\operatorname{Vol}(S) \leq(2 / 3) \operatorname{Vol}\left(V^{*}\right) \quad \text { and } \quad \Phi(S) \leq \frac{\phi^{3}}{19208 \ln ^{2}\left(\left|E^{*}\right| e^{4}\right)}
$$

Assuming such an $S$ exists, there is a CONGEST algorithm that finds a cut $C \subset V^{*}$ such that $\Phi(C) \leq 12 \phi$ in $O\left(D+\operatorname{poly}\left(\log \left|E^{*}\right|, 1 / \phi\right)\right)$ rounds, with failure probability $1 / \operatorname{poly}\left(\left|E^{*}\right|\right)$. Each vertex in $V^{*}$ knows whether or not it belongs to $C$.

### 2.2 Proof of Lemma 2.1

We prove Lemma 2.1 by presenting and analyzing a specific distributed algorithm, which makes use of the subroutines specified in Lemmas 2.3, 2.4, and 2.5.

Recall that we are given a subgraph with edge set $E^{\prime}$ and must ultimately return a partition of it into $E_{m}^{\prime} \cup E_{s}^{\prime} \cup E_{r}^{\prime}$. The algorithm initializes $E_{m}^{\prime} \leftarrow E^{\prime}, E_{s}^{\prime} \leftarrow \emptyset$, and $E_{r}^{\prime} \leftarrow \emptyset$. There are two types of special operations.

Remove. In an Remove operation, some edges are moved from $E_{m}^{\prime}$ to either $E_{s}^{\prime}$ or $E_{r}^{\prime}$. For the sake of a clearer presentation, each such operation is tagged Remove- $i$, for some index $i$.

Split. Throughout the algorithm we maintain a partition of the current set $E_{m}^{\prime}$. In a Split operation, the partition subdivided. Each such operation is tagged as Split-i, for some index $i$, such that Split-i occurs right after Remove- $i$.

Throughout the algorithm, we ensure that any part $E^{\star}$ of the partition of $E_{m}^{\prime}$ has an identifier that is known to all members of $V\left(E^{\star}\right)$. It is not required that each part forms a connected subgraph. The partition at the end of the algorithm, $E_{m}^{\prime}=\bigcup_{i=1}^{t} \mathcal{E}_{i}$, is the output partition.

Notations. Since we treat $E_{m}^{\prime}$ as the "active" edge set and $E_{s}^{\prime}$ and $E_{r}^{\prime}$ as repositories of removed edges, $\operatorname{deg}(v)$ refers to the degree of $v$ in the subgraph induced by the current $E_{m}^{\prime}$. We write $V_{\text {low }}=\left\{v \in V^{\prime} \mid \operatorname{deg}(v) \leq n^{\delta}\right\}$.

Algorithm. In the first step of the algorithm, move each edge $\{u, v\} \in E_{m}^{\prime}$ in the subgraph induced by $V_{\text {low }}$ to $E_{s, u}^{\prime}$, where $\operatorname{ID}(u)<\operatorname{ID}(v)$ (Remove-1). (Breaking ties by vertex-ID is critical to keep the orientation acyclic.)

After that, $E_{m}^{\prime}$ is divided into connected components. Assume these components are $G_{1}=$ $\left(V_{1}, E_{1}\right), G_{2}=\left(V_{2}, E_{2}\right), \ldots$, where $V_{i}=V\left(E_{i}\right)$. Let $D_{i}$ be the depth of a BFS tree rooted at an arbitrary vertex in $G_{i}$. In $O\left(D_{i}\right)$ rounds, the subgraph $G_{i}$ is assigned an identifier that is known to all vertices in $V_{i}$ (Split-1). Note that this step is done in parallel for each $G_{i}$, and the time for this step is different for each $G_{i}$. From now on there will be no communication between different subgraphs in $\left\{G_{1}, G_{2}, \ldots\right\}$, and we focus on one specific subgraph $G_{i}$ in the description of the algorithm.

Depending on how large $D_{i}$ is, there are two cases. If $D_{i} \geq 48 \log ^{2} m$, we go to Case 1 , otherwise we go to Case 2.

Case 1: In this case, we have $D_{i} \geq 48 \log ^{2} m$. Since there are no edges connecting two vertices in $V_{\text {low }}$, we can apply the High Diameter subroutine, Lemma 2.3, which finds a cut $(C, \bar{C})$ of $G_{i}$ such that $\min \left(|C|,\left|V_{i} \backslash C\right|\right) \geq \frac{D_{i}}{32} n^{\delta}$ and $\partial(C) \leq \min \left(\operatorname{Vol}(C), \operatorname{Vol}\left(V_{i} \backslash C\right)\right) /(12 \log m)$ in $O\left(D_{i}\right)$ rounds. Every vertex in $V_{i}$ knows whether it is in $C$ or not. All edges of the cut ( $C, \bar{C}$ ) are put into $E_{r}^{\prime}$ (Remove-2). Then $E_{i}$ splits into two parts according to the cut ( $C, \bar{C}$ ) (Split-2). After that, all vertices in $V_{i}$ terminate. (Observe that the part containing the BFS tree root is connected, but the other part is not necessarily connected.)

Case 2: In this case, we have $D_{i} \leq 48 \log ^{2} m$. Since $G_{i}=\left(V_{i}, E_{i}\right)$ is a small diameter graph, a vertex $v \in V_{i}$ is able broadcast a message to all vertices in $V_{i}$ very fast. We apply the Low Degree subroutine, Lemma 2.4, to obtain a partition $E_{i}=E^{\diamond} \cup E_{s}^{\diamond}$. We add all edges in $E_{s}^{\diamond}$ to $E_{s}^{\prime}$ in such a way that $E_{s, v}^{\prime} \leftarrow E_{s, v}^{\prime} \cup E_{s, v}^{\diamond}$ for all $v \in V_{i} \backslash V^{\diamond}$, where $V^{\diamond}=V\left(E^{\diamond}\right)$ (Remove-3).

After removing these edges, the remaining edges of $E_{i}$ are divided into several connected components, but all remaining vertices have degree larger than $n^{\delta} / 2$. Assume these connected components are $G_{i, 1}=\left(V_{i, 1}, E_{i, 1}\right), G_{i, 2}=\left(V_{i, 2}, E_{i, 2}\right), \ldots$ Let $D_{i, j}$ be the depth of the BFS tree from an arbitrary root vertex in $G_{i, j}$. In $O\left(D_{i, j}\right)$ rounds we compute such a BFS tree and assign an identifier that is known to all vertices in $V_{i, j}$ (Split-3). That is, the remaining edges in $E_{i}$ are partitioned into $E_{i, 1}, E_{1,2}, \ldots$

In what follows, we focus on one subgraph $G_{i, j}$ and proceed to Case 2-a or Case 2-b.
Case 2-a: In this case, $D_{i, j} \geq 48 \log ^{2} m$. The input specification of the High Diameter subroutine (Lemma 2.3) is satisfied, since every vertex has degree larger than $n^{\delta} / 2$. We apply the High Diameter subroutine to $G_{i, j}$. This takes $O\left(D_{i, j}\right)$ rounds. This case is similar to Case 1, and we do the same thing as what we do in Case 1, i.e., remove the edges in the cut found by the subroutine (Remove-4), split the remaining edges (Split-4), and then all vertices in $V_{i, j}$ terminate.

Case 2-b: In this case, $D_{i, j} \leq 48 \log ^{2} m$. Note that every vertex has degree larger than $n^{\delta} / 2$, and $G_{i, j}$ has small diameter. What we do in this case is to test whether $G_{i, j}$ has any low conductance cut; if yes, we will split $E_{i, j}$ into two components. To do so, we apply the Low Conductance subroutine, Lemma 2.5, with $\phi=\frac{1}{144 \log m}$. Based on the result, there are two cases.

Case 2-b-i: The subroutine finds a set of vertices $C$ that $\Phi(C) \leq 12 \phi=\frac{1}{12 \log m}$, and every vertex knows whether it is in $C$ or not. We move $\partial(C)$ to $E_{r}^{\prime}$ (Remove-5), and then split the remaining edges into two edge sets according to the cut $(C, \bar{C})$ (Split-5). After that, all vertices in $V_{i, j}$ terminate.

Case 2-b-ii: Otherwise, the subroutine does not return a subset $C$, and it means with probability at least $1-1 / \operatorname{poly}\left(\left|E_{i, j}\right|\right)=1-1 / \operatorname{poly}(n)$, there is no cut $(S, \bar{S})$ with conductance less than $\frac{\phi^{3}}{19208 \ln ^{2}\left(\left|E_{i, j}\right| e^{4}\right)}=\Theta\left(\log ^{-5} m\right)$. Recall the relation between the mixing time $\tau_{\text {mix }}\left(G_{i, j}\right)$ and the conductance $\Phi=\Phi_{G_{i, j}}: \Theta\left(\frac{1}{\Phi}\right) \leq \tau_{\text {mix }}\left(G_{i, j}\right) \leq \Theta\left(\frac{\log \left|V_{i, j}\right|}{\Phi^{2}}\right)[22]$. Therefore, w.h.p., $G_{i, j}$ has $O(\operatorname{poly} \log n)$ mixing time. All vertices in $V_{i, j}$ terminate without doing anything in this step.

Note that in the above calculation, we use the fact that every vertex in $V_{i, j}$ has degree larger than $n^{\delta} / 2$ in $G_{i, j}$, and this implies that $\left|V_{i, j}\right|=\Omega\left(n^{\delta}\right)$ and $\left|E_{i, j}\right|=\Omega\left(n^{2 \delta}\right)$, and so $\Theta(\log m)=$ $\Theta(\log n)=\Theta\left(\log \left|E_{i, j}\right|\right)=\Theta\left(\log \left|V_{i, j}\right|\right)$.

Analysis. We show that the output of $\mathcal{A}^{*}$ meets its specifications (C1)-(C6). Recall that $E_{m}^{\prime}=$ $\bigcup_{i=1}^{t} \mathcal{E}_{i}$ is the final partition of the edge set $E_{m}^{\prime}$ when all vertices terminate. Once an edge is moved from $E_{m}^{\prime}$ to either $E_{r}^{\prime}$ or $E_{s}^{\prime}$, it remains there for the rest of the computation. Condition (C1) follows from the fact that each time we do a split operation, the induced vertex set of each part is disjoint. Condition (C6) follows from the fact that each vertex knows which part of $E_{m}^{\prime}$ it belongs to after each split operation. In the rest of this section, we prove that the remaining conditions are met.

Claim 2.6. Condition (C2) is met.
Proof. Note that only Remove-1 and Remove-3 involve $E_{s}^{\prime}$. In Remove-1, any $E_{s, u}^{\prime}$ that becomes nonempty must have had $u \in V_{\text {low }}$, so $\operatorname{deg}(u) \leq n^{\delta}$ before Remove-1, and therefore $\left|E_{s, u}^{\prime}\right|+\operatorname{deg}(u) \leq n^{\delta}$ after Remove-1. In Remove-3, the Low Degree subroutine of Lemma 2.4 computes a partition $E_{i}=E^{\diamond} \cup E_{s}^{\diamond}$, and then we update $E_{s, u}^{\prime} \leftarrow E_{s, u}^{\prime} \cup E_{s, u}^{\diamond}$ for all $u \in V_{i} \backslash V^{\diamond}$. By Lemma 2.4, for any
$u$ such that $E_{s, u}^{\diamond} \neq \emptyset$, we have $\left|E_{s, u}^{\diamond}\right| \leq n^{\delta}$, and $u \notin V^{\diamond}$, where $V^{\diamond}$ is the vertex set induced by the remaining edge set $E^{\diamond}$. In other words, once $u$ puts at least one edge into $E_{s, u}^{\prime}$, we have $\operatorname{deg}(u)=0$ after Remove-3.

Claim 2.7. Conditions (C3) and (C4) are met.
Proof. We need to verify that in each part of the algorithm, we either spend only $O(\operatorname{poly} \log n)$ rounds, or the size of the current component shrinks by $\tilde{\Omega}\left(n^{\delta}\right)$ vertices per round.

After removing all edges in the subgraph induced by $V_{\text {low }}$, the rest of $E^{\prime}$ is partitioned into connected components $\mathcal{E}_{1}, \mathcal{E}_{2}, \ldots$ Consider one such component $\mathcal{E}_{i}$, and suppose it goes to Case 1. We find a sparse cut $(C, \bar{C})$, and moving $\partial(C)$ to $E_{r}^{\prime}$ breaks $\mathcal{E}_{i}$ into $\mathcal{E}_{i}^{1}$ and $\mathcal{E}_{i}^{2}$. By Lemma 2.3, we have $\min (|C|,|\bar{C}|) \geq \frac{D_{i}}{32} n^{\delta}$, so the size of both $V\left(\mathcal{E}_{i}^{1}\right)=C$ and $V\left(\mathcal{E}_{i}^{2}\right)=\bar{C}$ are at most $\left|V\left(\mathcal{E}_{i}\right)\right|-$ $\frac{D_{i}}{32} n^{\delta} \leq n^{\prime}-\Omega\left(D_{i}\right) n^{\delta}$. Since the running time for each vertex in $V\left(\mathcal{E}_{i}^{1}\right)$ and $V\left(\mathcal{E}_{i}^{2}\right)$ is $O\left(D_{i}\right)$, the condition (C3-2) is met.

Now suppose that $\mathcal{E}_{i}$ goes to Case 2. Note that the total time spent before it reaches Case 2 is $O\left(D_{i}\right)=$ poly $\log n$. In Case 2 we execute the Low Degree subroutine of Lemma 2.4, and let the time spent in this subroutine be $\tau$. By Lemma 2.4, it is either the case that (i) $\tau=O\left(D_{i}\right)$ or (ii) the remaining vertex set $V^{\diamond}$ satisfies $\left|V\left(E_{i}\right)\right|-\left|V^{\diamond}\right|=\Omega\left(\tau n^{\delta}\right)$. In other words, if we spend too much time (i.e., $\omega\left(D_{i}\right)$ ) on this subroutine, we must lose $\Omega\left(n^{\delta}\right)$ vertices per round.

After that, $\mathcal{E}_{i}$ is split into $\mathcal{E}_{i, 1}, \mathcal{E}_{i, 2}, \ldots$ We consider the set $\mathcal{E}_{i, j}$. If $\mathcal{E}_{i, j}$ goes to Case 2 -a, then the analysis is the same as that in Case 1, and so (C3-2) is met.

Now suppose that $\mathcal{E}_{i, j}$ goes to Case $2-\mathrm{b}$. Note that the time spent during the Low Conductance subroutine of Lemma 2.5 is $O(\operatorname{poly} \log n)$. Suppose that a low conductance cut $(C, \bar{C})$ is found (Case 2-b-i). Since the cut has conductance less than $\frac{1}{12 \log m}$, by the fact that every vertex has degree higher than $n^{\delta} / 2$, we must have $\min (|C|,|\bar{C}|)=\Omega\left(n^{\delta}\right)$. Assume $\mathcal{E}_{i, j} \backslash \partial(C)$ is split into $\mathcal{E}_{i, j}^{1}$ and $\mathcal{E}_{i, j}^{2}$. The size of both $V\left(\mathcal{E}_{i, j}^{1}\right)$ and $V\left(\mathcal{E}_{i, j}^{2}\right)$ must be at most $\left|V\left(\mathcal{E}_{i, j}\right)\right|-\Omega\left(n^{\delta}\right)$. Thus, (C3-2) holds for both parts $\mathcal{E}_{i, j}^{1}$ and $\mathcal{E}_{i, j}^{2}$.

Suppose that no cut $(C, \bar{C})$ is found (Case 2-b-ii). If the running time $K$ among vertices in $V_{i, j}$ is $O$ (poly $\log n$ ), then (C3-1) holds. Otherwise, we must have $\left|V_{i, j}\right| \leq n^{\prime}-\tilde{\Omega}\left(K n^{\delta}\right)$ due to the Low Degree subroutine, and so (C3-2) holds.

Condition (C4) follows from the the above proof of (C3), since for each part of the algorithm, it is either the case that (i) this part takes $O$ (poly $\log n)$ time, or (ii) the number of vertices in the current subgraph is reduced by $\tilde{\Omega}\left(n^{\delta}\right)$ per round.

Claim 2.8. Condition (C5) is met.
Proof. Condition (C5) says that after the algorithm $\mathcal{A}^{*}$ completes, $\left|E_{r}^{\prime}\right| \leq f$, where

$$
f=\left(\left|E^{\prime}\right| \log \left|E^{\prime}\right|-\sum_{i=1}^{t}\left|\mathcal{E}_{i}\right| \log \left|\mathcal{E}_{i}\right|\right) /(6 \log m)
$$

We prove the stronger claim that this inequality holds at all times w.r.t. the current edge partition $\mathcal{E}_{1} \cup \cdots \cup \mathcal{E}_{t}$ of $E_{m}^{\prime}$. In the base case this is clearly true, since $t=1$ and $E^{\prime}=E_{m}^{\prime}=\mathcal{E}_{1}$ and $E_{r}^{\prime}=\emptyset$. Moving edges from $E_{m}^{\prime}$ to $E_{s}^{\prime}$ increases $f$ and has no effect on $E_{r}^{\prime}$, so we only have to consider the movement of edges from $E_{m}^{\prime}$ to $E_{r}^{\prime}$. Note that this only occurs in Remove- $i$ and Split- $i$, for $i \in\{2,4,5\}$, where in these operations we find a cut $(C, \bar{C})$ and split one of the parts $\mathcal{E}_{j}$ according to the cut. In all cases we have

$$
|\partial(C)| \leq \frac{\min (\operatorname{Vol}(C), \operatorname{Vol}(\bar{C}))}{12 \log m}
$$

Suppose that removing $\partial(C)$ splits $\mathcal{E}_{j}$ into $\mathcal{E}_{j}^{1}$ and $\mathcal{E}_{j}^{2}$, with $\left|\mathcal{E}_{j}^{1}\right| \leq\left|\mathcal{E}_{j}^{2}\right|$ and $C=V\left(\mathcal{E}_{j}^{1}\right)$. We bound the change in $\left|E_{r}^{\prime}\right|$ and $f$ separately. Clearly

$$
\Delta\left|E_{r}^{\prime}\right|=|\partial(C)| \leq \frac{2\left|\mathcal{E}_{j}^{1}\right|+\partial(C)}{12 \log m} \leq \frac{\left|\mathcal{E}_{j}^{1}\right|}{6 \log m}+\frac{\partial(C)}{12 \log m} .
$$

and

$$
\begin{aligned}
\Delta f & =\frac{1}{6 \log m} \cdot\left(\left|\mathcal{E}_{j}\right| \log \left|\mathcal{E}_{j}\right|-\sum_{k \in\{1,2\}}\left|\mathcal{E}_{j}^{k}\right| \log \left|\mathcal{E}_{j}^{k}\right|\right) \\
& \geq \frac{1}{6 \log m} \cdot\left(\left|\mathcal{E}_{j}^{1}\right| \log \left(\left|\mathcal{E}_{j}\right| /\left|\mathcal{E}_{j}^{1}\right|\right)+\partial(C) \log \left|\mathcal{E}_{j}\right|\right) \\
& \left.>\Delta\left|E_{r}^{\prime}\right| \quad \quad \text { (Because }\left|\mathcal{E}_{j}^{1}\right|<\left|\mathcal{E}_{j}\right| / 2 .\right)
\end{aligned}
$$

Thus, $\left|E_{r}^{\prime}\right| \leq f$ also holds after Remove- $i$ and Split- $i$, for $i \in\{2,4,5\}$.

## 3 Algorithm for Finding a Sparse Cut

Recall the in our decomposition routine, we search for a sparse cut in a subgraph $G^{*}=\left(V\left(E^{*}\right), E^{*}\right)$ of $G$. To simplify notation, we use $n=\left|V\left(G^{*}\right)\right|$ and $m=\left|E\left(G^{*}\right)\right|$ to be the number of vertices and edges in the subgraph. In this section we prove Lemma 2.5, which concerns an efficient distributed analogue of Spielman and Teng's [39,40] Nibble routine.

Many existing works $[3,8,27,39]$ have shown that looking at the distribution of random walks is a good approach to finding a sparse cut. The basic idea is to first sample a source vertex $s$ according to the degree distribution, i.e., the probability that $v$ is sampled is $\operatorname{deg}(v) /(2 m)$, and do a lazy random walk from $s$. Assume there is a sparse cut $S$ with conductance $\Phi(S)$, and $\operatorname{Vol}(S) \leq \operatorname{Vol}(V) / 2$. If $s \in S$, then the probability distribution of the random walk will be mostly confined to $S$ within the initial $t_{0}=O\left(\frac{1}{\Phi(S)}\right)$ steps. A common way to utilize this observation is to sort the vertices $\left(v_{1}, \ldots, v_{n}\right)$ in decreasing order of their random walk probability, and it is guaranteed that for some choice of $j$, the subset $C=\left\{v_{1}, \ldots, v_{j}\right\}$ is a sparse cut that is approximately as good as $S$.

The papers [27,38] adapted this approach to the CONGEST model. If the cut $S$ satisfies that $b \cdot 2|E| \leq \operatorname{Vol}(S)$ (i.e., $S$ has balance $b$ ), then a cut $C$ satisfying $\Phi(C)=O(\sqrt{\Phi(S) \log n})$ can be found in $\tilde{O}(D+1 /(b \Phi(S)))$ rounds. The algorithm is inefficient when $1 / b=\Theta(|E| / \operatorname{Vol}(S))$ is large. The main source of this inefficiency is that if we sample a vertex $s$ according to the degree distribution, then the probability that $s \in S$ is only $O(b)$. This implies that we have to calculate many random walk distributions before we find a desired sparse cut. If we calculate these random walk distributions simultaneously, then we may suffer from a huge congestion issue.

Spielman and Teng [39] show that a random walk distribution with truncation (rounding a probability to zero when it becomes too small) can reveal a sparse cut, provided the starting vertex of the random walk is good. The main contribution of this section is a proof that the Spielman-Teng method for finding cuts of conductance roughly $\phi$ can be implemented in poly $\left(\phi^{-1}, \log n\right)$ time in the CONGEST model, i.e., with no dependence on the balance parameter $b$.

Terminology. We first review some definitions and results from Spielman and Teng [39]. Let $A$ be the adjacency matrix of the graph $G$. We assume 1-1 correspondance between $V$ and $\{1, \ldots, n\}$. In a lazy random walk, the walk stays at the current vertex with probability $1 / 2$ and otherwise moves to a random neighbor of the current vertex. The matrix realizing this walk can be expressed
as $T=\left(A D^{-1}+I\right) / 2$, where $D$ is the diagonal matrix with $(d(1), \ldots, d(n))$ on the diagonal, and $d(i)=\operatorname{deg}(i)$.

Let $p_{t}^{v}$ be the probability distribution of the lazy random walk that begins at $v$ and walks for $t$ steps. In the limit, as $t \rightarrow \infty, p_{t}(x)$ approaches $d(x) / 2 m$, so it is natural to measure $p_{t}(x)$ relative to this baseline.

$$
\rho_{t}(x)=p_{t}(x) / d(x)
$$

Define $\pi_{t}$ to be the permutation that sorts $V=\{1, \ldots, n\}$ in decreasing order of $\rho_{t}$-values, breaking ties by vertex ID. (We never actually compute $\pi_{t}$. To implement our algorithms, it suffices that given $\rho_{t}(u), \rho_{t}(v), \operatorname{ID}(u), \operatorname{ID}(v)$, we can determine whether or not $u$ precedes $v$ according to $\pi_{t}$.)

$$
\rho_{t}\left(\pi_{t}(i)\right) \geq \rho_{t}\left(\pi_{t}(i+1)\right), \text { for all } i
$$

Let $p$ be a distribution on $V$. The truncation operation $[p]_{\epsilon}$ rounds $p(x)$ to zero if it falls below a threshold that depends on $x$.

$$
[p]_{\epsilon}(x)= \begin{cases}p(x) & \text { if } p(x) \geq 2 \epsilon d(x) \\ 0 & \text { otherwise }\end{cases}
$$

The truncated random walk starting at vertex $v$ is defined as follows. In subsequent discussion we may omit $v$ if it is known implicitly.

$$
\begin{aligned}
\tilde{p}_{0}^{v}(x) & = \begin{cases}1 & x=v \text { and } 1 \geq 2 \epsilon d(x) \\
0 & \text { otherwise }\end{cases} \\
\tilde{p}_{t}^{v} & =\left[T \tilde{p}_{t-1}\right]_{\epsilon}
\end{aligned}
$$

The description of the algorithm Nibble and Lemma 3.1 in [39] implies the following lemma. ${ }^{4}$
Lemma 3.1 ( [39] ). For each $\phi \leq 1$, define the parameters

$$
\begin{aligned}
t_{0} & =\frac{49 \ln \left(m e^{4}\right)}{\phi^{2}} \\
\text { and } \gamma & =\frac{5 \phi}{392 \ln \left(m e^{4}\right)} .
\end{aligned}
$$

For each subset $S \subset V$ satisfying

$$
\begin{aligned}
\operatorname{Vol}(S) & \leq \frac{2}{3} \cdot \operatorname{Vol}(V) \\
\text { and } \Phi(S) & \leq \frac{\phi^{3}}{19208 \ln ^{2}\left(m e^{4}\right)}
\end{aligned}
$$

there exists a subset $S^{g} \subseteq S$ with the following properties. First, $\operatorname{Vol}\left(S^{g}\right) \geq \operatorname{Vol}(S) / 2$. Second, $S^{g}$ is partitioned into $S^{g}=\bigcup_{b=1}^{\log m} S_{b}^{g}$ such that if a random walk is initiated at any $v \in S_{b}^{g}$ with truncation parameter $\epsilon=\frac{\phi}{56 \ln \left(m e^{4}\right) t_{0} 2^{b}}$, then there exists a number $t \in\left[1, t_{0}\right]$ and an index $j$ such that the following four conditions are met for the cut $\left.C=\left\{\tilde{\pi}_{t}^{v}(1), \ldots, \tilde{\pi}_{t}^{v}(j)\right)\right\}$.
(i) $\Phi(C) \leq \phi$,

[^3](ii) $\tilde{\rho}_{t}\left(\tilde{\pi}_{t}(j)\right) \geq \gamma / \operatorname{Vol}(C)$,
(iii) $\operatorname{Vol}(C \cap S) \geq(4 / 7) 2^{b-1}$,
(iv) $\operatorname{Vol}(C) \leq(5 / 6) \operatorname{Vol}(V)$.

In subsequent discussion, with respect to a given parameter $\phi \leq 1$, for any subset $S \subset V$ satisfying the condition of Lemma 3.1, we fix a subset $S^{g} \subseteq S$ and its decomposition $S^{g}=\bigcup_{b=1}^{\log m} S_{b}^{g}$ to be any choices satisfying Lemma 3.1.

### 3.1 Distributed Algorithm

Now we give our algorithm Distributed Nibble. To simplify things, we present it as a sequential algorithm, and prove in Lemma 3.5 that it can be implemented efficiently in the CONGEST model. For any permutation $\pi$, we use the notation $\pi(i . . j)$ to denote the set $\{\pi(i), \pi(i+1), \ldots, \pi(j)\}$.

```
Algorithm 1 Distributed Nibble
    Input: \(\phi\).
    for parameter \(b=1\) to \(\lceil\log m\rceil\) do
        Set parameters \(t_{0}=49 \ln \left(m e^{4}\right) / \phi^{2}\), and \(\epsilon_{b}=\frac{\phi}{\left.56 \ln \left(m e^{4}\right) t_{0} 2^{b}\right)}\), as in Lemma 3.1.
        (1) Independently randomly sample \(K=c \log m \cdot \frac{\operatorname{Vol}(V)}{2^{b}}\) vertices \(v_{1}, \ldots, v_{K}\) proportional to
        their degrees, where \(c\) is a large enough constant.
        Initialize \(\tilde{p}_{0}^{v_{i}}\).
        for \(t=1\) to \(t_{0}\), for every \(v_{i}\) do
            (2) calculate \(\tilde{p}_{t}^{v_{i}}=\left[T \tilde{p}_{t-1}^{v_{i}}\right]_{\epsilon_{b}}\)
            Denote \(j_{\text {max }}\) as the largest index such that \(\tilde{p}_{t}^{v_{i}}\left(\tilde{\pi}_{t}^{v_{i}}\left(j_{\max }\right)\right)>0\).
            for \(x=0\) to \(\log _{1+\phi}(5 / 6) \operatorname{Vol}(V)\) do
                (3) Set \(j \leq j_{\text {max }}\) to be the largest index that \(\operatorname{Vol}\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right) \leq(1+\phi)^{x}\).
                    (4) If \(\Phi\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right) \leq 12 \phi\), output the sparse cut \(C=\tilde{\pi}_{t}^{v_{i}}(1 . . j)\) and halt.
            end for
        end for
    end for
    Return failed.
```

From Lemma 3.1 we know that we can obtain a cut $C$ with some good properties if we start the truncated random walk at a vertex $v \in S_{b}^{g}$ with parameter $\epsilon_{b}$. Therefore, what we do in Distributed Nibble is to just sample sufficiently many vertices as the starting points of random walks so that with sufficiently high probability at least one them is in the set $S_{b}^{g}$. The danger here is that calculating all these random walk distributions simultaneously may be infeasible if any part of the graph becomes too congested.

In this section we analyze the behavior of Distributed Nibble (as a sequential algorithm) and prove that it operates correctly. In Section 3.2 we argue that Distributed Nibble can be implemented efficiently in the CONGEST model, in poly $(\log m, 1 / \phi)$ time.

Roughly speaking, Lemma 3.2 shows that if the sets $\pi(1 . . j)$ and $\pi\left(1 . . j^{\prime}\right)$ have similar volume, then the cuts resulting from these two sets have similar sparsity. This justifies lines (3) and (4) of Distributed Nibble and allows us to examine a small number of prefixes of the permutation $\tilde{\pi}_{t}^{v_{i}}$.

Lemma 3.2. Let $\pi$ be any permutation, and let $\phi \leq 1 / 12$. If, for some index $j, \Phi(\pi(1 . . j)) \leq \phi$ and $\operatorname{Vol}(\pi(1 . . j)) \leq(5 / 6) \operatorname{Vol}(V)$, then $\Phi\left(\pi\left(1 . . j^{\prime}\right)\right) \leq 12 \phi$ for all indices $j^{\prime}>j$ such that

$$
\operatorname{Vol}\left(\pi\left(1 . . j^{\prime}\right)\right) \leq(1+\phi) \operatorname{Vol}(\pi(1 . . j))
$$

Proof. Let $x=\operatorname{Vol}(\pi(1 . . j))$ and $y=\operatorname{Vol}\left(\pi\left(1 . . j^{\prime}\right)\right)$. Recall that $2 m=2|E|=\operatorname{Vol}(V)$, and so $x \leq(5 / 6) \operatorname{Vol}(V)=(5 / 6) 2 m$. We have $x \leq y \leq(1+\phi) x$. Since $x \leq(5 / 6) 2 m$ and $\phi \leq 1 / 12$, we have $\phi x \leq x / 12 \leq(2 m-x) / 2$. Therefore,

$$
2 m-y \geq 2 m-x-\phi x \geq(2 m-x) / 2
$$

We calculate an upper bound of $\Phi\left(\pi\left(1 . . j^{\prime}\right)\right)$ as follows.

$$
\Phi\left(\pi\left(1 . . j^{\prime}\right)\right)=\frac{\partial\left(\pi\left(1 . . j^{\prime}\right)\right)}{\min (y, 2 m-y)} \leq \frac{\partial(\pi(1 . . j))+\sum_{i=j+1}^{j^{\prime}} d(\pi(i))}{\min (x,(2 m-x) / 2)} \leq \frac{\partial(\pi(1 . . j))+\phi x}{\min (x,(2 m-x)) / 2} \leq 12 \phi
$$

We explain the details of the derivation. The first inequality is due to $x \leq y$ and $(2 m-x) / 2 \leq$ $2 m-y$, which follow from the above discussion. The second inequality is due to the fact that $\sum_{i=j+1}^{j^{\prime}} d(\pi(i))=\operatorname{Vol}\left(\pi\left(1 . . j^{\prime}\right)\right)-\operatorname{Vol}(\pi(1 . . j)) \leq \phi \cdot \operatorname{Vol}(\pi(1 . . j))=\phi x$. For the third inequality, note that $\frac{\partial(\pi(1 . . j))}{\min (x,(2 m-x))} \leq \phi$ and $\frac{\phi x}{\min (x,(2 m-x))} \leq 5 \phi$, since $x \leq(5 / 6) 2 m$.

Lemma 3.3. Let $S \subset V$ be any subset satisfying

$$
\operatorname{Vol}(S) \leq(2 / 3) \operatorname{Vol}(V) \quad \text { and } \quad \Phi(S) \leq \frac{\phi^{3}}{19208 \ln ^{2}\left(m e^{4}\right)}
$$

Then there exists a number $b$ such that $\operatorname{Vol}\left(S_{b}^{g}\right) \geq 2^{b} / 32$.
Proof. Denote $x=\operatorname{Vol}(S)$. From Condition (iii) of Lemma 3.1 we deduce that if $S_{b}^{g} \neq \emptyset$, then there exists a set of vertices $C$ such that $\operatorname{Vol}(S) \geq \operatorname{Vol}(C \cap S) \geq(4 / 7) 2^{b-1}$. Thus, for all $b$ such that $b \geq\lceil\log x\rceil+2$, we must have $S_{b}^{g}=\emptyset$. If the statement of this lemma is false, i.e., $\operatorname{Vol}\left(S_{b}^{g}\right)<2^{b} / 32$ for all $b$, then

$$
\operatorname{Vol}\left(S^{g}\right) \leq \sum_{b=1}^{\lceil\log x\rceil+1} \frac{2^{b}}{32}<\frac{2^{\lceil\log x\rceil+2}}{32}<x / 4
$$

which contradicts the requirement $\operatorname{Vol}\left(S^{g}\right) \geq \operatorname{Vol}(S) / 2$ specified in Lemma 3.1.
Lemma 3.4 (Correctness). For any $\phi \leq 1 / 12$, if there exists a subset $S \subset V$ satisfying

$$
\operatorname{Vol}(S) \leq(2 / 3) \operatorname{Vol}(V) \quad \text { and } \quad \Phi(S) \leq \frac{\phi^{3}}{19208 \ln ^{2}\left(m e^{4}\right)}
$$

then Distributed Nibble outputs a set of vertices $C$ such that $\Phi(C) \leq 12 \phi$ with probability at least $1-1 / \operatorname{poly}(m)$.

Proof. From Lemma 3.3 we know there exists a number $b$ such that $\operatorname{Vol}\left(S_{b}^{g}\right) \geq 2^{b} / 32$. Since we sample $v_{i}$ proportional to the degree distribution,

$$
\operatorname{Pr}\left[v_{i} \in S_{b}^{g}\right]=\frac{\operatorname{Vol}\left(S_{b}^{g}\right)}{\operatorname{Vol}(V)} \geq \frac{2^{b}}{32 \cdot \operatorname{Vol}(V)} .
$$

Since we sample $K=c \log m \cdot \frac{\operatorname{Vol}(V)}{2^{b}}$ number of vertices,

$$
\operatorname{Pr}\left[\exists i \text { s.t. } v_{i} \in S_{b}^{g}\right] \geq 1-\left(1-\frac{2^{b}}{32 \operatorname{Vol}(V)}\right)^{c \log m \frac{\mathrm{Vol}(V)}{2^{b}}} \geq 1-m^{-\Omega(c)} .
$$

Now we focus on the truncated random walk starting at this vertex $v_{i} \in S_{b}^{g}$. We fix two numbers $t \in\left[1, t_{0}\right]$ and $j$ such that the four conditions in Lemma 3.1 are satisfied. In particular, Condition (i) and Condition (iv) in Lemma 3.1 say that

$$
\begin{aligned}
\operatorname{Vol}\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right) & \leq(5 / 6) \operatorname{Vol}(V), \\
\Phi\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right) & \leq \phi .
\end{aligned}
$$

Therefore, we are able to apply Lemma 3.2 , and so we have $\Phi\left(\tilde{\pi}_{t}^{v_{i}}\left(1 . . j^{\prime}\right)\right) \leq 12 \phi$ for all indices $j^{\prime}$ such that $\operatorname{Vol}\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right) \leq \operatorname{Vol}\left(\tilde{\pi}_{t}^{v_{i}}\left(1 . . j^{\prime}\right)\right) \leq(1+\phi) \operatorname{Vol}\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right)$.

In Distributed Nibble, we search for a cut with target volume $(1+\phi)^{x}$, for all possible integers $x$. Note that Condition (ii) in Lemma 3.1 implies $j \leq j_{\text {max }}$. Therefore, in Step (3) of Distributed Nibble, at least one index $j^{\star}$ picked by the algorithm satisfies

$$
\operatorname{Vol}\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right) \leq \operatorname{Vol}\left(\tilde{\pi}_{t}^{v_{i}}\left(1 . . j^{\star}\right)\right) \leq(1+\phi) \operatorname{Vol}\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right) .
$$

By Lemma 3.2, the cut $C=\tilde{\pi}_{t}^{v_{i}}\left(\left\{1, \ldots, j^{\star}\right\}\right)$ associated with this index $j^{\star}$ found in Step (4) meets the requirement $\Phi(C) \leq 12 \phi$ of the lemma.

### 3.2 Implementation

We show how to implement Distributed Nibble in the CONGEST model. The goal of this section is to prove Lemma 3.5. Note that Lemma 2.5 is a consequence of Lemmas 3.4 and 3.5.
Lemma 3.5. Distributed Nibble can be implemented in the CONGEST model using $O(D+$ $\log ^{9} m / \phi^{10}$ ) rounds, with success probability $1-1 / \operatorname{poly}(m)$, where $D$ is the diameter of graph. If Distributed Nibble outputs a set C successfully, then each vertex knows whether or not it belongs to $C$.

To prove this lemma, we shall analyze Distributed Nibble step by step.
Lemma 3.6 (Step (1)). The samples for every level b (from 1 to $\lceil\log m\rceil$ ) can be generated in $O(D+\log m)$ time.

Proof. We build a BFS tree rooted at an arbitrary vertex $x$. For each vertex $v$, define $s(v)$ as the sum of $d(u)$ for each $u$ in the subtree rooted at $v$. In $O(D)$ rounds we can let each vertex $v$ learn the number $s(v)$ by a bottom-up traversal of the BFS tree.

In the beginning, for each $b=1, \ldots,\lceil\log m\rceil$, we generate $K_{b}=c \log m \frac{\operatorname{Vol}(V)}{2^{b}}$ number of $b$-tokens at the root $x$. Let $L=\Theta(D)$ be the number of layers in the BFS tree. For $i=1, \ldots, L$, the vertices of layer $i$ do the following. When a $b$-token arrives at $v$, the token disappears at $v$ with probability $d(v) / s(v)$ and $v$ includes itself in the $b$ th sample; otherwise, $v$ sends the token to a child $u$ with probability $\frac{s(u)}{s(v)-d(v)}$. Note that $v$ only needs to tell each child $u$ how many $b$-tokens $u$ gets. Thus, for each $b$, the process of choosing $K_{b}=c \log m \frac{\operatorname{Vol}(V)}{2^{b}}$ vertices from the degree distribution can be done in $L$ rounds. By pipelining, we can do this for all $b$ in $O(D+\log m)$ rounds.

This method has the virtue of selecting exactly $K_{b}$ vertices in the $b$ th sample. We can also select $K_{b}$ vertices in expectation, in just $O(D)$ time, simply by computing $\operatorname{Vol}(V)$ with a BFS tree, disseminating it to all vertices, and letting each $v$ join the sample independently with probability $K_{b} \operatorname{deg}(v) / \operatorname{Vol}(V)$.

It is not obvious why Step (2) of Distributed Nibble should be efficiently implementable in the CONGEST model. Before analyzing it, we give some helpful lemmas about lazy random walks.

Lemma 3.7 ([39] ). For all $u, v$, and $t, \rho_{t}^{v}(u)=\rho_{t}^{u}(v)$.
Proof. This lemma was observed in [39] without proof. For the sake of completeness, we provide a short proof here. A sequence of vertices $W=\left(x_{0}, x_{1}, \ldots, x_{t}\right)$ is called a walk of length $t$ if $x_{i+1} \in N\left(x_{i}\right) \cup\left\{x_{i}\right\}$ for each $i \in[0, t)$. We write $\operatorname{Pr}[W]$ to be the probability that the first $t$ steps of a lazy random walk starting at $x_{0}$ tracks $W$. Let $W^{R}=\left(x_{t}, x_{t-1}, \ldots, x_{0}\right)$ be the reversal of $W$.

Let $\mathcal{W}_{t}^{u, v}$ be the set of walks of length $t$ starting at $u$ and ending at $v$. It is clear that $\rho_{t}^{u}(v)=$ $\sum_{W \in \mathcal{W}_{t}^{u, v}} \operatorname{Pr}[W] / d(v)$ and $\rho_{t}^{v}(u)=\sum_{W \in \mathcal{W}_{t}^{v, u}} \operatorname{Pr}[W] / d(u)$. Since $\mathcal{W}_{t}^{v, u}=\left\{W^{R} \mid W \in \mathcal{W}_{t}^{u, v}\right\}$, to prove the lemma it suffices to show that $\operatorname{Pr}[W] / d(v)=\operatorname{Pr}\left[W^{R}\right] / d(u)$ for each $W \in \mathcal{W}_{t}^{u, v}$.

Fix any $W \in \mathcal{W}_{t}^{u, v}$ and let $W_{*}=\left(y_{0}, \ldots, y_{s}\right)$ be the subsequence of $W$ resulting from splicing out immediate repetitions in $W$. It is clear that $\operatorname{Pr}[W]=2^{-t} \cdot \prod_{i=0}^{s-1} 1 / d\left(y_{i}\right)$, and so

$$
\frac{\operatorname{Pr}[W]}{d(v)}=\frac{\operatorname{Pr}[W]}{d\left(y_{s}\right)}=2^{-t} \cdot \prod_{i=0}^{s} \frac{1}{d\left(y_{i}\right)}=\frac{\operatorname{Pr}\left[W^{R}\right]}{d\left(y_{0}\right)}=\frac{\operatorname{Pr}[W]}{d(u)} .
$$

Lemma 3.8. Fix the parameter $b$ (which influences $\epsilon_{b}$ and hence the truncation operation of the random walk) and define

$$
Z_{t}(u)=\left\{v_{i} \mid v_{i} \text { is in the bth sample and } \tilde{p}_{t-1}^{v_{i}}(u)>0\right\} .
$$

For every vertex $u$ and every $t$, with high probability, $\left|Z_{t}(u)\right| \leq O\left(\log ^{3} m / \phi^{3}\right)$.
Proof. Define $S=\left\{v \in V \mid \tilde{p}_{t-1}^{v}(u)>0\right\}$. By definition $Z_{t}(u)=S \cap\left\{v_{1}, \ldots, v_{K_{b}}\right\}$. For each $v \in S$, we have $p_{t-1}^{v}(u) \geq \tilde{p}_{t-1}^{v}(u) \geq 2 \epsilon_{b} d(u)$. Recall that $p_{t-1}$ is the probability distribution obtained after $t-1$ steps of the lazy random walk without truncation. By Lemma 3.7,

$$
p_{t-1}^{u}(v)=\left(p_{t-1}^{v}(u) / d(u)\right) d(v) \geq 2 \epsilon_{b} \cdot d(v) .
$$

Therefore, $2 \epsilon_{b} \cdot \operatorname{Vol}(S) \leq \sum_{v \in S} p_{t-1}^{u}(v) \leq 1$, and so $\operatorname{Vol}(S) \leq \frac{1}{2 \epsilon_{b}}$, which implies

$$
\operatorname{Pr}\left[v_{i} \in S\right] \leq \frac{1}{2 \epsilon_{b} \cdot \operatorname{Vol}(V)}
$$

Recall that $t_{0}=\frac{49 \ln \left(m e^{4}\right)}{\phi^{2}}$ and $\epsilon_{b}=\frac{\phi}{56 \ln \left(m e^{4}\right) t_{0} 2^{b}}$. Rewrite the number $K_{b}=c \log m \frac{\mathrm{Vol}(V)}{2^{b}}$ as $K_{b}=\Theta\left(\epsilon_{b} \cdot \operatorname{Vol}(V) \cdot \log ^{3} m / \phi^{3}\right)$. Since each of $v_{1}, \ldots, v_{K_{b}}$ is chosen independently, using a Chernoff bound we conclude that there exists a constant $c^{\prime}>0$ depending on $c$ such that

$$
\operatorname{Pr}\left[\left|Z_{t}(u)\right|>c^{\prime} \log ^{3} m / \phi^{3}\right] \leq \exp \left(-\Omega\left(\log ^{3} m / \phi^{3}\right)\right)
$$

Lemma 3.9 (Step (2)). Fix the parameter b. Suppose each vertex $v$ knows $\tilde{p}_{t-1}^{v_{i}}(v)$, for all $v_{i}$ in the bth sample. Then with high probability, each vertex $v$ can calculate $\tilde{p}_{t}^{v_{i}}(v)$, for all $v_{i}$, within $O\left(\log ^{3} \mathrm{~m} / \phi^{3}\right)$ rounds.

Proof. The normal way to calculate $\left[T \tilde{p}_{t-1}(u)\right]_{\epsilon_{b}}$ is as follows. For each $v_{i}$, each vertex $v$ broadcasts the number $\frac{\tilde{p}_{t-1}^{v_{i}}(v)}{2 d(v)}$ to all its neighbors, and then $v$ collects messages from neighbors. The vertex $v$ can calculate $\tilde{p}_{t}^{v_{i}}(v)$ locally by adding $\tilde{p}_{t-1}^{v_{i}}(v) / 2$ and all numbers received from its neighbors, then applying the truncation operation. Note that a straightforward analysis of this protocol leads to a terrible round complexity, since we have to do this for each $v_{i}$.

One crucial observation is that a vertex $v$ does not need to care about those $v_{i}$ with $\tilde{p}_{t-1}^{v_{i}}(v)=0$ at time $t$. We modify this protocol a little bit in such a way that we never send a number if it is 0 . Define $Z_{t}(u)=\left\{v_{i} \mid \tilde{p}_{t-1}^{v_{i}}(u)>0\right\}$ as in Lemma 3.8, and so each vertex $v$ only needs to spend $\left|Z_{t}(v)\right|$ rounds to simulate the time step $t$ of the lazy random walk. By Lemma 3.8 and the discussion above, we have proved that Step (2) can be executed in $O\left(\log ^{3} m / \phi^{3}\right)$ time, for every $v_{i}$ and any specific $t$.

Lemma 3.10 (Steps (3,4)). Fix parameters $b$, $t$ and $x$. Steps (3) and (4) can be implemented in $O\left(\log ^{6} m / \phi^{7}\right)$ rounds for all $v_{i}$ in the bth sample. For any sparse cut $C$ found in Step (4), every vertex in $C$ knows that it belongs to $C$.

Proof. Now we focus on the random walk starting at $v_{i}$. Let $U=\left\{u \mid \exists t^{\prime} \leq t, \tilde{p}_{t^{\prime}}^{v_{i}}(u)>0\right\}$. $U$ must be a connected vertex set. Obviously all $\tilde{\pi}_{t}^{v_{i}}(j)$ for $j \leq j_{\max }$ are in $U$. We build a BFS tree of $U$ rooted at $v_{i}$, which has $t+1$ levels. We will execute Step (3) and Step (4) by sending requests from the root to all vertices in $U$, collecting information from $U$ to the root, and making a decision locally at the root. Recall that each $v_{i}$ has its own BFS tree, and in general a vertex $u$ belongs to multiple BFS trees for different $v_{i}$. Luckily, each vertex $u$ only belongs to the BFS tree of those $v_{i} \in \bigcup_{1 \leq t \leq t_{0}} Z_{t}(u)$, so with only a $t_{0} \cdot \max _{u, t}\left|Z_{t}(u)\right|=O\left(\log ^{4} m / \phi^{5}\right)$ overhead of running time, we can do Step (3) and Step (4) for all $v_{i}$ in parallel.

To find each index $j$ specified in Step (3), we can do a "random binary search" on vertices in $U$. Let $\tilde{\pi}=\tilde{\pi}_{t}^{v_{i}}$ and $\tilde{\rho}=\tilde{\rho}_{t}^{v_{i}}$ be with respect to $\tilde{p}_{t}^{v_{i}}$. Note that by our choice of $U$ we can assume $U$ is a prefix set of $\tilde{\pi}$. We maintain two indices $L$ and $R$ that control the search space. Initially, $L \leftarrow 1$ and $R \leftarrow|U|$. In each iteration, we randomly pick one vertex $\tilde{\pi}(j)$ among $\tilde{\pi}(L . . R)$ and calculate $\operatorname{Vol}(\tilde{\pi}(1, \cdots, j))$ by broadcasting $\tilde{\rho}(\tilde{\pi}(j))$ to all vertices in $U$ and propagating information up the BFS tree. If $\operatorname{Vol}(\tilde{\pi}(1, \cdots, j)) \leq(1+\phi)^{x}$, we update $L \leftarrow j$; otherwise we let $R=j-1$. In each iteration, with probability $1 / 2$ we sample $j$ in the middle half of $[L, R]$ and the size of search space $[L, R]$ shrinks by a factor of at least $3 / 4$. Therefore, w.h.p., after $O(\log m)$ iterations, we will have isolated $L=R=j$. Each iteration can be done in $O(t)=O\left(t_{0}\right)$ rounds. Due to the congestion overhead, Step (3) can be implemented in $O\left(\log m \cdot t_{0} \cdot \log ^{4} m / \phi^{5}\right)=O\left(\log ^{6} m / \phi^{7}\right)$ rounds.

Step (4) can be done by simply collecting information about $\partial\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right)$ and $\operatorname{Vol}\left(\tilde{\pi}_{t}^{v_{i}}(1 . . j)\right)$; its round complexity is of a lower order than that of Step (3). If the root $v_{i}$ finds a cut $C$ with $\Phi(C) \leq 12 \phi$, it broadcasts $\tilde{\rho}(\tilde{\pi}(j))$ to all vertices in $U$ to let the vertices in $C$ know that they are in $C$. Note that for each vertex $u$ in $U$, it can infer whether it is in $C$ by comparing $\tilde{\rho}(u)$ and $\tilde{\rho}(\tilde{\pi}(j))$.

Proof of Lemma 3.5. Combining Lemmas 3.6, 3.9, and 3.10, the running time in Step (1) is $O(D+$ $\log m)$, Step (2) is $O\left(\log ^{5} m / \phi^{5}\right)$, and Steps (3) and (4) are $O\left(\log ^{9} m / \phi^{10}\right)$. The dominating term $O\left(\log ^{9} m / \phi^{10}\right)$ comes from enumerating $\log m \cdot t_{0} \cdot \log m / \phi=\Theta\left(\log ^{3} m / \phi^{3}\right)$ combinations of $(b, t, x)$, spending $O\left(\log ^{6} m / \phi^{7}\right)$ rounds for each combination.

Whenever a vertex $v_{i}$ finds a sparse cut $C$, it broadcasts a message to the entire graph saying that it has found a cut, and this takes $O(D)$ rounds. If multiple cuts are found by different vertices, we can select exactly one cut, breaking ties arbitrarily. A more opportunistic version of the algorithm could also take a maximal independent set of compatible cuts.

## 4 Triangle Enumeration

We use the routing algorithm from [16,17]. Theorem 2 was first stated in [16, Theorem 1.2] with round complexity $\tau_{\text {mix }}(G) \cdot 2^{O(\sqrt{\log n \log \log n})}$; this was recently improved to $\tau_{\text {mix }}(G) \cdot 2^{O(\sqrt{\log n})}$ in [17].

Theorem 2 ( $[16,17]$ ). Consider a graph $G=(V, E)$ and a set of point-to-point routing requests, each given by the IDs of the corresponding source-destination pair. If each vertex $v$ is the source and the destination of at most $\operatorname{deg}(v) \cdot 2^{O(\sqrt{\log n})}$ messages, there is a randomized distributed algorithm that delivers all messages in $\tau_{\text {mix }}(G) \cdot 2^{O(\sqrt{\log n})}$ rounds, w.h.p., in the CONGEST model.

Remark 1. The claim of Theorem 2 appears to be unproven for arbitrary ID-assignments (HsinHao Su, personal communication, 2018), but is true for well-behaved ID-assignments, which we illustrate can be computed efficiently in CONGEST. In [16, 17] each vertex $v \in V$ simulates $\operatorname{deg}(v)$ virtual vertices in a random graph $G_{0}$ which is negligibly close to one drawn from the Erdős-Rényi distribution $\mathcal{G}(2 m, p)$ for some $p$. Presumably the IDs of $v$ 's virtual vertices are $(\operatorname{ID}(v), 1), \ldots,(\operatorname{ID}(v), \operatorname{deg}(v))$. It is proven [16, 17] that effecting a set of routing requests in $G_{0}$ takes $2^{O(\sqrt{\log n})}$ time in $G_{0}$; however, to translate a routing request $\operatorname{ID}(x) \rightsquigarrow \operatorname{ID}(y)$ in $G$ to $G_{0}$, it seems necessary to map it (probabilistically) to $(\operatorname{ID}(x), i) \rightsquigarrow(\operatorname{ID}(y), j)$, where $i, j$ are chosen uniformly at random from $[1, \operatorname{deg}(x)]$ and $[1, \operatorname{deg}(y)]$, respectively. (This is important for the global congestion guarantee that $y$ 's virtual nodes receive roughly equal numbers of messages from all sources.) This seems to require that $x$ know how to compute $\operatorname{deg}(y)$ or an approximation thereof based on $\operatorname{ID}(y)$. Arbitrary ID-assignments obviously do not betray this information.

Lemma 4.1. In $O(D+\log n)$ time we can compute an ID assignment ID : $V \rightarrow\{1, \ldots,|V|\}$ and other information such that $\operatorname{ID}(u)<\operatorname{ID}(v)$ implies $\lfloor\log \operatorname{deg}(u)\rfloor \leq\lfloor\log \operatorname{deg}(v)\rfloor$, and any vertex $u$ can locally compute $\lfloor\log \operatorname{deg}(v)\rfloor$ for any $v$.

Proof. Build a BFS tree from an arbitrary vertex $x$ in $O(D)$ time. In a bottom-up fashion, each vertex in the BFS tree calculates the number of vertices $v$ in its subtree having $\lfloor\log \operatorname{deg}(v)\rfloor=i$, for $i=0, \ldots, \log n$. This takes $O(D+\log n)$ time by pipelining. At this point the root $x$ has the counts $n_{0}, \ldots, n_{\log n}$ for each degree class, where $n=\sum_{i} n_{i}$. It partitions up the ID-space so that all vertices in class- 0 get IDs from [ $1, n_{0}$ ], class- 1 from $\left[n_{0}+1, n_{0}+n_{1}\right]$, and so on. The root broadcasts the numbers $n_{0}, \ldots, n_{\log n}$, and disseminates the IDs to all nodes according to their degrees. (In particular, the root gives each child $\log n$ intervals of the ID-space, which they further subdivide, sending $\log n$ intervals to the grandchildren, etc.) With pipelining this takes another $O(D+\log n)$ time. Clearly knowing $n_{0}, \ldots, n_{\log n}$ and $\operatorname{ID}(v)$ suffice to calculate $\lfloor\log \operatorname{deg}(v)\rfloor$.

Lemma 4.1 gives us a good ID-assignment to apply Theorem 2. It is also useful in our triangle enumeration application. Roughly speaking, vertices with larger degrees also have more bandwidth in the CONGEST model, and therefore should be responsible for learning about larger subgraphs and enumerating more triangles.

Before we present our triangle enumeration algorithm for general graphs, we address the important special case of finding triangles with at least one edge in a component of high conductance.

### 4.1 Triangle Enumeration in High Conductance Graphs

Recall that our graph decomposition routine returns a tripartition $E_{m} \cup E_{s} \cup E_{r}$. Triangles that intersect $E_{s}$ will be enumerated separately. The purpose of this section is to provide a routine to enumerate triangles that intersect $E_{m}$, i.e., they are completely contained in $E_{m}$ or have one edge in $E_{m}$ and two in $E_{r}$. Whereas each component of $E_{m}$ has low mixing time, we can say nothing about the mixing time of $E_{m}$ plus all incident $E_{r}$ edges.

Definitions. The input is a subgraph $G_{\text {in }}=\left(V_{\text {in }}, E_{\text {in }}\right)$ with low mixing time, together with some edges $E_{\text {out }}$ joining vertices in $V_{\text {in }}$ to $V-V_{\text {in }}$. Let $\operatorname{deg}_{\text {in }}(v)$ and $\operatorname{deg}_{\text {out }}(v)$ be the number of $E_{\text {in }}$ and $E_{\text {out }}$ edges incident to $v$. In this section $n=\left|V_{\text {in }}\right|$ and $m=\left|E_{\text {in }}\right|$.

Theorem 3. Suppose that $G_{\mathrm{in}}$ and $E_{\text {out }}$ meet the following conditions:
(i) For each $v \in V_{\text {in }}, \operatorname{deg}_{\text {in }}(v) \geq \operatorname{deg}_{\text {out }}(v)$.
(ii) $\tau_{\text {mix }}\left(G_{\text {in }}\right)=n^{o(1)}$.

In the CONGEST model, all triangles in $E_{\text {in }} \cup E_{\text {out }}$ can be counted and enumerated, w.h.p., in $O\left(n^{1 / 3+o(1)}\right)$ rounds.

Note that Theorem 3 applies to the class of graphs with $n^{o(1)}$ mixing time, letting $E_{\text {out }}=\emptyset$. We first describe the algorithm behind Theorem 3 and then analyze it in Lemmas 4.2-4.5.

The Easy Case. We first check whether any vertex $v^{\star} \in V\left(E_{\text {in }} \cup E_{\text {out }}\right)$ has $\operatorname{deg}_{\text {in }}\left(v^{\star}\right)+\operatorname{deg}_{\text {out }}\left(v^{\star}\right) \geq$ $m /\left(20 n^{1 / 3} \log n\right)=\zeta$. If so, we apply Theorem 2 to the subgraph $G_{\text {in }}^{+}$induced by $V_{\text {in }} \cup\left\{v^{\star}\right\}$ and have every vertex $u \in V_{\text {in }}$ transmit to $v^{\star}$ all its incident edges in $E_{\text {in }} \cup E_{\text {out. }}{ }^{5}$ Condition (i) of Theorem 3 implies that $\left|E_{\text {in }} \cup E_{\text {out }}\right| \leq 3 m$, so the total volume of messages entering $v^{\star}$ is $O(m)$. Therefore the routing takes $O\left(\tau_{\text {mix }}\left(G_{\text {in }}^{+}\right) \cdot 2^{O(\sqrt{\log n})} \cdot m / \zeta\right)=n^{1 / 3+o(1)}$, and thereafter, $v^{\star}$ can report all triangles in $E_{\text {in }} \cup E_{\text {out }}$. In the analysis of the following steps, we may assume that the maximum degree in the graph induced by $E_{\text {in }} \cup E_{\text {out }}$ is at most $m /\left(20 n^{1 / 3} \log n\right)$.

Vertex Classes. Let $\delta=2^{\lfloor\log (2 m / n)\rfloor}$ be the average degree in $G_{\text {in }}$, rounded down to the nearest power of 2 . Write $\operatorname{deg}_{\text {in }}(v)=k_{v} \cdot \delta$, and call $v$ a class-0 vertex if $k_{v} \in[0,1 / 2)$ and a class-i vertex if $k_{v} \in\left[2^{i-2}, 2^{i-1}\right)$. We use the fact that

$$
\sum_{v \in V_{\text {in }}: k_{v} \geq 1 / 2} 2 k_{v} \geq n
$$

By applying Lemma 4.1 to reassign IDs, we may assume that the ID-space of $V_{\text {in }}$ is $\left\{1, \ldots,\left|V_{\text {in }}\right|\right\}$ and that any vertex can compute the class of $v$, given $\operatorname{ID}(v)$.

Randomized Partition. Our algorithm is a randomized adaptation of the CONGESTED-CLIQUE algorithm of [9]. We partition the set $V\left(E_{\text {in }} \cup E_{\text {out }}\right)$ into $V_{1} \cup \cdots \cup V_{n^{1 / 3}}$ locally, without communication. Each vertex $v \in V\left(E_{\text {in }} \cup E_{\text {out }}\right)$ selects an integer $r_{v} \in\left[1, n^{1 / 3}\right]$ uniformly at random, joins $V_{r v}$, and transmits ' $r_{v}$ ' to its immediate neighbors. We allocate the (less than) $n$ triads

$$
\mathfrak{T}=\left\{\left(j_{1}, j_{2}, j_{3}\right) \mid 1 \leq j_{1} \leq j_{2} \leq j_{3} \leq n^{1 / 3}\right\}
$$

to the vertices in $V_{\text {in }}$ in the following way. Enumerate the vertices in increasing order of ID. If $v$ is class- 0 , then skip $v$. If $v$ is class- $i, i \geq 1$, then $k_{v}<2^{i-1} / \delta$. Allocate to $v$ the next $2^{i} / \delta \geq 2 k_{v}$ triads from $\mathcal{T}$, and stop whenever all triads are allocated. By Lemma 4.1, every vertex can calculate the class of any other, and can therefore perform this allocation locally, without communication. A vertex $v \in V$ that is assigned a triad $\left(j_{1}, j_{2}, j_{3}\right)$ is responsible of learning the set of all edges $E\left(V_{j_{1}}, V_{j_{2}}\right) \cup E\left(V_{j_{2}}, V_{j_{3}}\right) \cup E\left(V_{j_{1}}, V_{j_{3}}\right)$ and reporting/counting those triangles ( $x_{1}, x_{2}, x_{3}$ ) with $x_{k} \in V_{j_{k}} .{ }^{6}$

[^4]Transmitting Edges. Every vertex knows the IDs of its neighbors and which part of the vertex partition they are in. For each $v \in V_{\text {in }}$, each incident edge $(v, u) \in E_{\text {in }} \cup E_{\text {out }}$, and each index $r^{*} \in\left[1, n^{1 / 3}\right], v$ transmits the message " $(v, u), r_{v}, r_{u}$ " to the unique vertex $x$ handling the triad on $\left\{r_{u}, r_{v}, r^{*}\right\}$. Observe that the total message volume is exactly $\Theta\left(m n^{1 / 3}\right)$.

We analyze the behavior of this algorithm in the CONGEST model, where the last step is implemented by applying Theorem 2 to $G_{\text {in }}$. Recall from Condition (i) of Theorem 3 that the number of edges in the graph we consider, $\bar{m}=\left|E_{\text {in }} \cup E_{\text {out }}\right|$, is in the range $[m, 3 m]$.

Lemma 4.2. Consider a graph with $\bar{m}$ edges and $\bar{n}$ vertices. We generate a subset $S$ by letting each vertex join $S$ independently with probability $p$. Suppose that the maximum degree is $\Delta \leq \bar{m} p / 20 \log \bar{n}$ and $p^{2} \bar{m} \geq 400 \log ^{2} \bar{n}$. Then, w.h.p., the number of edges in the subgraph induced by $S$ is at most $6 p^{2} \bar{m}$.

Proof. For an edge $e_{i}$, define $x_{i}=1$ if both two endpoints of edge $e_{i}$ join $S$, otherwise $x_{i}=0$. Then $X=\sum_{i=1}^{\bar{m}} x_{i}$ is the number of edges in the subgraph induced by $S$. We have $\mathbb{E}[X]=p^{2} \bar{m}$, and by Markov's inequality,

$$
\operatorname{Pr}[X \geq 6 \mathbb{E}[X]]=\operatorname{Pr}\left[X^{c} \geq(6 \mathbb{E}[X])^{c}\right] \leq \frac{1}{6^{c}} \frac{\mathbb{E}\left[X^{c}\right]}{p^{2 c} \bar{m}^{c}},
$$

where $c=5 \log \bar{n}$ is a parameter.

$$
\begin{aligned}
\mathbb{E}\left[X^{c}\right] & =\sum_{i_{1}, \ldots, i_{c} \in[1, \bar{m}]} \mathbb{E}\left[\prod_{j=1}^{c} x_{i_{j}}\right] \\
& =\sum_{k=2}^{2 c} f_{k} \cdot p^{k},
\end{aligned}
$$

where $f_{k}$ is the number of choices $\left\{i_{1}, \ldots, i_{c} \in[1, \bar{m}]\right\}$ such that the number of distinct endpoints in the edge set $e_{i_{1}}, \ldots, e_{i_{c}}$ is $k$.

For any choice of $\left(i_{1}, \ldots, i_{c} \in[1, \bar{m}]\right)$, we project it to a vector $\left\langle k_{1}, \ldots, k_{c}\right\rangle \in\{0,1,2\}^{c}$, where $k_{j}$ indicates the number of endpoints of $e_{i_{j}}$ that overlap with the endpoints of the edges $e_{i_{1}}, \ldots, e_{i_{j-1}}$. Note that $2 c-\sum k_{j}$ is the number of distinct endpoints in the edge set $\left\{e_{i_{1}}, \ldots, e_{i_{c}}\right\}$. We fix a vector $\left\langle k_{1}, \ldots, k_{c}\right\rangle$ and count how many choices of $\left(i_{1}, \ldots, i_{c}\right)$ project to this vector.

Suppose that the edges $e_{i_{1}}, \ldots, e_{i_{j-1}}$ are fixed. We bound the number of choices of $e_{i_{j}}$ as follows. If $k_{j}=0$, the number of choices is clearly at most $m$. If $k_{j}=1$, the number of choices is at most $(2 c)(p \bar{m} / 20 \log \bar{n})$, since one of its endpoints (which overlaps with the endpoints of the edges $e_{i_{1}}, \ldots, e_{i_{j-1}}$ ) has at most $2 c$ choices, and the other endpoint (which does not overlap with the endpoints of the edges $e_{i_{1}}, \ldots, e_{i_{j-1}}$ ) has at most $\Delta \leq \bar{m} p / 20 \log \bar{n}$ choices. If $k_{j}=2$, the number of choices is at most $(2 c)^{2}$.

Based on the above calculation, we upper bound $f_{k}$ as follows. In the calculation, $x$ is the number of indices $j$ such that $k_{j}=1$, and $y$ is the number of indices $j$ that $k_{j}=2$. Note that $\binom{c}{x}\binom{c-x}{y}$ is the number of distinct vectors $\left\langle k_{1}, \ldots, k_{c}\right\rangle$ realizing the given parameters $c, x$, and $y$.

$$
\begin{aligned}
f_{k} & \leq \sum_{x+y \leq c, 2 c-x-2 y=k} \bar{m}^{c-x-y}\binom{c}{x}\binom{c-x}{y}\left(\frac{2 c p \bar{m}}{20 \log \bar{n}}\right)^{x}\left(4 c^{2}\right)^{y} \\
& \leq \sum_{x+y \leq c, 2 c-x-2 y=k} \bar{m}^{c} 3^{c}\left(\frac{2 c p}{20 \log \bar{n}}\right)^{x}\left(\frac{4 c^{2}}{\bar{m}}\right)^{y}
\end{aligned}
$$

$$
\begin{aligned}
& \leq \sum_{x+y \leq c, 2 c-x-2 y=k}(3 \bar{m})^{c}\left(\frac{2 c p}{20 \log \bar{n}}\right)^{x+2 y} \\
& \leq c(3 \bar{m})^{c}\left(\frac{2 c p}{20 \log \bar{n}}\right)^{2 c-k}
\end{aligned}
$$

The third inequality is due to the fact $p^{2} \bar{m} \geq 400 \log ^{2} \bar{n}$, which implies $(2 c p / 20 \log \bar{n})^{2} \geq\left(4 c^{2} / \bar{m}\right)$. Using the fact that $\frac{2 c}{20 \log \bar{n}} \leq 1 / 2$, we upper bound $\mathbb{E}\left[X^{c}\right]$ as follows.

$$
\begin{aligned}
\mathbb{E}\left[X^{c}\right] & \leq \sum_{k=2}^{2 c} f_{k} \cdot p^{k} \\
& =c(3 \bar{m})^{c} p^{2 c} \sum_{k=2}^{2 c}\left(\frac{2 c}{20 \log \bar{n}}\right)^{2 c-k} \\
& <2 c(3 \bar{m})^{c} p^{2 c} .
\end{aligned}
$$

Therefore,

$$
\operatorname{Pr}[X \geq 6 \mathbb{E}[X]] \leq \frac{1}{6^{c}} \frac{\mathbb{E}\left[X^{c}\right]}{p^{2 c} \bar{m}^{c}} \leq \frac{2 c 3^{c}}{6^{c}} \leq \frac{10 \log \bar{n}}{\bar{n}^{5}}
$$

Note that the probability can be amplified to $\bar{n}^{-t}$ for any constant $t$ by setting $c=t \log \bar{n}$ and using different constants in the statement of the lemma.

Lemma 4.3. For all $j_{1}, j_{2} \in\left[1, n^{1 / 3}\right]$, with high probability, $\left|E\left(V_{j_{1}}, V_{j_{2}}\right)\right| \leq 24 m / n^{2 / 3}$.
Proof. For the case $j_{1}=j_{2}$, we apply Lemma 4.2 to the subgraph on $E_{\text {in }} \cup E_{\text {out }}$ having $\bar{m}$ edges and $\bar{n}$ vertices, with sampling probability $p=n^{-1 / 3}$. One may verify that the maximum degree is at most $m p /(20 \log n)<\bar{m} p /(20 \log \bar{n})$ and that $p^{2} \bar{m} \geq n^{1 / 3} \geq 400 \log ^{2} \bar{n}$. By Lemma 4.2, we conclude that $\operatorname{Pr}\left[\left|E\left(V_{j_{1}}\right)\right|>6 \bar{m} / n^{2 / 3}\right] \leq \frac{10 \log n}{n^{5}}$.

For the case $j_{1} \neq j_{2}$, we can apply the same analysis to the subgraph induced by $V_{j_{1}} \cup V_{j_{2}}$ with $p=2 n^{-1 / 3}$ and conclude that $\operatorname{Pr}\left[\left|E\left(V_{j_{1}}, V_{j_{2}}\right)\right|>24 \bar{m} / n^{2 / 3}\right] \leq \frac{10 \log n}{n^{5}}$. By a union bound over all $n^{2 / 3}$ choices of $j_{1}$ and $j_{2}$, the stated upper bound holds everywhere, w.h.p.

Lemma 4.4. With high probability, each vertex $v \in V_{\text {in }}$ receives $O\left(\operatorname{deg}_{\text {in }}(v) \cdot n^{1 / 3}\right)$ edges.
Proof. Consider any vertex $v \in V_{\text {in }}$. If $k_{v}<1 / 2$, then $v$ receives no message; otherwise $v$ is responsible for between $2 k_{v}$ and $4 k_{v}$ triads, and $v$ collects the edge set $E\left(V_{j_{1}}, V_{j_{2}}\right)$ for at most $12 k_{v}$ pairs of $V_{j_{1}}, V_{j_{2}}$. By Lemma 4.3, w.h.p., $\left|E\left(V_{j_{1}}, V_{j_{2}}\right)\right|=O\left(\mathrm{~m} / \mathrm{n}^{2 / 3}\right)$. Hence $v$ receives

$$
O\left(m / n^{2 / 3}\right) \cdot 12 k_{v}=O\left(\operatorname{deg}_{\mathrm{in}}(v) \cdot n^{1 / 3}\right)
$$

messages, with high probability.
Lemma 4.5. Each vertex $v \in V_{\text {in }}$ sends $O\left(\operatorname{deg}_{\text {in }}(v) \cdot n^{1 / 3}\right)$ edges with probability 1.
Proof. By Condition (i) of Theorem 3, $v \in V_{\text {in }}$ is responsible for $\operatorname{deg}_{\text {in }}(v)+\operatorname{deg}_{\text {out }}(v) \leq 2 \operatorname{deg}_{\text {in }}(v)$ incident edges, and each is involved in exactly $n^{1 / 3}$ triads.

Lemmas $4.2-4.5$ show that the message volume into/out of every vertex is close to its expectation. By applying Theorem 2 and Lemma 4.1, all messages can be routed in $O\left(n^{1 / 3+o(1)}\right)$ time. This concludes the proof of Theorem 3.
Corollary 1. Let $G$ be a graph with $\tau_{\text {mix }}(G)=n^{o(1)}$. In the CONGEST model, Triangle Detection, Enumeration, and Counting can be solved on $G$, with high probability, in $n^{1 / 3+o(1)}$ time.

### 4.2 Triangle Enumeration and Counting in General Graphs

The algorithm for Theorem 4 is based on an $n^{1 / 2}$-decomposition. Since the connected components induced by $E_{m}$ have low mixing time, we can solve Triangle Enumeration/Counting on them very efficiently using Theorem 3, in $n^{1 / 3+o(1)}$ time, i.e., much less than the time required to compute the $n^{1 / 2}$-decomposition.

Theorem 4. In the CONGEST model, Triangle Detection, Counting, and Enumeration can be solved, w.h.p., in $\tilde{O}\left(n^{1 / 2}\right)$ rounds.

Proof. The underlying graph is $G=(V, E)$. We set the parameter $\delta=1 / 2$. By Theorem 1, we compute an $n^{\delta}$-decomposition $E=E_{m} \cup E_{s} \cup E_{r}$ using $\tilde{O}\left(n^{1-\delta}\right)$ rounds. We divide the task of enumerating triangles into three cases. By ensuring that every triangle is output by exactly one vertex, this algorithm also solves Triangle Counting.

Case 1: All Triangles Intersecting $E_{s}$. We handle this case as follows. By Condition (b) of Definition 2, $E_{s}=\bigcup_{v \in V} E_{s, v}$, where $\left\{E_{s, v}\right\}$ defines an acyclic $n^{\delta}$-orientation. We let each $v$ announce $E_{s, v}$ to all its neighbors, in $O\left(n^{\delta}\right)$ time. For the Triangle Counting application it is important that every triangle $\{x, y, z\}$ intersecting $E_{s}$ be reported by exactly one vertex. If $(x, y)$ and $(x, z)$ are oriented and $\operatorname{ID}(y)<\operatorname{ID}(z)$, then $y$ detects and reports the triangle. If $(x, z)$ is oriented, $\{x, y\}$ is unoriented, and $\{y, z\}$ is unoriented or oriented as $(z, y)$, then $y$ detects and reports the triangle. If $(x, z),(y, z)$ are oriented but $\{x, y\}$ is not, and $\operatorname{ID}(y)<\operatorname{ID}(x), y$ reports the triangle. ${ }^{7}$

Case 2: Some Triangles Intersecting $E_{m}$. Consider a single connected component $G_{\mathrm{in}}=$ $\left(V_{\text {in }}, E_{\text {in }}\right)$ induced by $E_{m}$, which has mixing time $n^{o(1)}$. We classify vertices in $V_{\text {in }}$ as good or bad depending on whether they naturally satisfy Condition (i) of Theorem 3. A vertex is good if $\operatorname{deg}_{\mathrm{in}}(v) \geq \operatorname{deg}_{E_{r}}(v)$. Let $E_{\text {out }}$ be the subset of $E_{r}$-edges incident to good vertices in $V_{\mathrm{in}}$, and let $E_{r}^{\text {new }}$ be the subset of $E_{m}$-edges incident to bad vertices in $V_{\text {in }}$. We now apply Theorem 3 to enumerate/count all triangles in the edge set $E_{\text {in }} \cup E_{\text {out }}$. (Because triangles contained in $E_{r}^{\text {new }}$ will also be found in Case 3, the Triangle Counting algorithm should refrain from including these in the tally for Case 2.)

Case 3: Triangles Contained in $E_{r}^{\text {new }} \cup E_{r}$. Let $E_{r}^{\text {new }}$ denote the union of all such sets, over all components in $E_{m}$. Since each edge in $E_{r}^{\text {new }}$ can be charged to an endpoint of an edge in $E_{r}$, $\left|E_{r}^{\text {new }} \cup E_{r}\right| \leq 3\left|E_{r}\right|=|E| / 2$. We apply the algorithm recursively to the graph induced by $E_{r}^{\text {new }} \cup E_{r}$. The depth of the recursion is obviously at $\operatorname{most} \log _{2} m$.

Round Complexity. Computing an $n^{\delta}$-decomposition $E=E_{m} \cup E_{s} \cup E_{r}$ takes $\tilde{O}\left(n^{1-\delta}\right)$ rounds. The algorithm for Case 1 takes $O\left(n^{\delta}\right)$ rounds. The algorithm for Case 2 takes $O\left(n^{1 / 3+o(1)}\right)$ rounds. The number of recursive calls (Case 3 ) is $\log m$. Thus, the overall round complexity is

$$
\log m \cdot\left(O\left(n^{\delta}\right)+\tilde{O}\left(n^{1-\delta}\right)+O\left(n^{1 / 3+o(1)}\right)\right)=\tilde{O}\left(n^{1 / 2}\right)
$$

[^5]
### 4.3 Subgraph Enumeration

In this section we show that Corollary 1 can be extended to enumerating $s$-vertex subgraphs in $O\left(n^{(s-2) / s+o(1)}\right)$ rounds. Note that the $\Omega\left(n^{1 / 3} / \log n\right)$ lower bound for triangle enumeration on Erdős-Rényi graphs $\mathcal{G}(n, 1 / 2)$ [21] can be generalized to an $\Omega\left(n^{(s-2) / s} / \log n\right)$ lower bound for enumerating $s$-vertex subgraphs. This implies that Theorem 5 is nearly optimal on $\mathcal{G}(n, 1 / 2)$.

Theorem 5. Let $s=O(1)$ be any constant. Given a graph $G$ of $n$ vertices with $\tau_{\text {mix }}(G)=n^{o(1)}$, we can list all s-vertex subgraphs of $G$ in $O\left(n^{(s-2) / s+o(1)}\right)$ rounds, w.h.p., in the CONGEST model.

This theorem is proved using a variant of Lemma 3, also with $\mathcal{G}=G$ and $\left|N_{\mathcal{G}}(v)\right|=\operatorname{deg}(v)$. The proof of Theorem 5 is almost the same as that of Lemma 3, and so in what follows we only highlight the difference.

Similarly, we assume the maximum degree is at most $m /\left(20 n^{(s-2) / s} \log n\right)$, since otherwise we can apply Theorem 2 to let $v$ learn the entire set $E$ in $O\left(n^{(s-2) / s+o(1)}\right)$ rounds, and we are done after that.

We partition $V$ into $n^{1 / s}$ subsets $V_{1}, \ldots, V_{n^{1 / s}}$. Instead of considering triads, here we consider s-tuples: $\left\{\left(i_{1}, \ldots, i_{s}\right) \mid 1 \leq i_{1} \leq \ldots \leq i_{s} \leq n^{1 / s}\right\}$. After a vertex $v$ learns the entire edge set $\cup_{j_{1}, j_{2} \in[1, s]} E\left(V_{i_{j_{1}}}, V_{i_{j_{2}}}\right)$, it has ability to list all $s$-vertex subgraphs in which the $j$ th vertex is in $V_{i_{j}}$. We prove a variant of Lemma 4.3, as follows.

Lemma 4.6. W.h.p., $\left|E\left(V_{i}, V_{j}\right)\right| \leq 24 m / n^{2 / s}$ for all $i, j \in\left[1, n^{1 / s}\right]$.
Proof. We set $p=n^{-1 / s}$. The maximum degree is at most $m /\left(20 n^{(s-2) / s} \log n\right) \leq m p / 20 \log n$, and $p^{2} m \geq m / n^{2 / s} \geq 400 \log ^{2} n$. By applying Lemma 4.2 and use the same analysis in Lemma 4.3, we conclude this lemma.

Proof of Theorem 5. Here we only consider the time complexity to delivery all messages. Consider a vertex $v$. If $k_{v}<1 / 2$, then $v$ receives no message. Otherwise $v$ is responsible for between $2 k_{v}$ and $4 k_{v}$ tuples, and $v$ collects $E\left(V_{i}, V_{j}\right)$ for at most $4 s^{2} k_{v}$ pairs $\left(V_{i}, V_{j}\right)$. By Lemma 4.6, w.h.p., $\left|E\left(V_{i}, V_{j}\right)\right| \leq 24 m / n^{2 / s}$ for all $i, j$. Hence the number of edges $v$ received is at most $24 m / n^{2 / s} \cdot 4 s^{2} k_{v}=\bar{O}\left(\operatorname{deg}(v) \cdot n^{(s-2) / s}\right)$.

Note that each vertex $v$ sends at most $O\left(\operatorname{deg}(v) n^{(s-2) / s}\right)$ messages since for each incident edge $e$ of $v$, there are at most $O\left(s^{2} n^{(s-2) / s}\right)$ tuples involving $e$. By Theorem 2, the delivery of all messages can be done in $O\left(n^{(s-2) / s+o(1)}\right)$ rounds, w.h.p.

## 5 Conclusion

In this paper we have shown that all variants of Triangle Detection, Enumeration, and Counting can be solved in $\tilde{O}\left(n^{1 / 2}\right)$ rounds in the CONGEST model. Moreover, we have shown that with better distributed graph partitioning technology, our algorithm can be implemented in $O\left(n^{1 / 3+o(1)}\right)$ rounds, nearly achieving the $\Omega\left(n^{1 / 3} / \log n\right)$ CONGEST lower bound of Izumi and LeGall [21]. ${ }^{8}$ Thus, the main question left open by our work is whether such an ideal graph decomposition can be computed efficiently, in $O\left(n^{1 / 3+o(1)}\right)$ time or even $O(\operatorname{polylog}(n))$ time. We conjecture that this is, in fact, possible.

[^6]Conjecture 1. In the CONGEST model, the edge set $E$ can be partitioned into $E=E_{m} \cup E_{r}$ in polylog( $n$ ) time such that $\left|E_{r}\right|<|E| / 2$ and the components induced by $E_{m}$ have conductance $\Omega(1 / \operatorname{polylog}(n))$ and hence $O(\operatorname{polylog}(n))$ mixing time.

If Conjecture 1 is true, then the upper bound for triangle enumeration, detection, and counting is roughly

$$
\min \left\{\tilde{O}(\Delta), n^{1 / 3+o(1)}\right\}
$$

i.e., depending on the magnitude of $\Delta$, we should execute one of two algorithms. For which edge-densities and which problems (counting, detection, enumeration) can this upper bound be improved? Is there a third algorithm that is substantially superior to these two in some regime?

## References

[1] A. Abboud, K. Censor-Hillel, S. Khoury, and C. Lenzen. Fooling views: A new lower bound technique for distributed computations under congestion. arXiv preprint arXiv:1711.01623, 2017.
[2] U. Agarwal, V. Ramachandran, V. King, and M. Pontecorvi. A deterministic distributed algorithm for exact weighted all-pairs shortest paths in $\tilde{O}\left(n^{3 / 2}\right)$ rounds. In Proceedings $38 t h$ ACM Symposium on Principles of Distributed Computing (PODC), pages ??-??, 2018.
[3] R. Andersen, F. R. K. Chung, and K. J. Lang. Local partitioning for directed graphs using PageRank. Internet Mathematics, 5(1):3-22, 2008.
[4] R. Andersen, S. O. Gharan, Y. Peres, and L. Trevisan. Almost optimal local graph clustering using evolving sets. J. ACM, 63(2):15:1-15:31, May 2016.
[5] S. Arora, B. Barak, and D. Steurer. Subexponential algorithms for unique games and related problems. J. ACM, 62(5):42:1-42:25, Nov. 2015.
[6] K. Censor-Hillel, P. Kaski, J. H. Korhonen, C. Lenzen, A. Paz, and J. Suomela. Algebraic methods in the congested clique. Distributed Computing, 2016.
[7] A. Czumaj and C. Konrad. Detecting cliques in CONGEST networks. arXiv preprint arXiv:1807.01070, 2018.
[8] A. Das Sarma, S. Gollapudi, and R. Panigrahy. Sparse cut projections in graph streams. In Proceedings $1^{77}$ th European Symposium on Algorithms (ESA), pages 480-491, 2009.
[9] D. Dolev, C. Lenzen, and S. Peled. "Tri, tri again": Finding triangles and small subgraphs in a distributed setting. In Proceedings 26th International Symposium on Distributed Computing (DISC), pages 195-209, 2012.
[10] A. Drucker, F. Kuhn, and R. Oshman. On the power of the congested clique model. In Proceedings 33rd ACM Symposium on Principles of Distributed Computing (PODC), pages 367-376, 2014.
[11] M. Elkin. Distributed exact shortest paths in sublinear time. In Proceedings 49th Annual ACM Symposium on Theory of Computing (STOC), pages 757-770, 2017.
[12] M. Elkin. A simple deterministic distributed MST algorithm, with near-optimal time and message complexities. In Proceedings 37th ACM Symposium on Principles of Distributed Computing (PODC), pages 157-163, 2017.
[13] G. Even, O. Fischer, P. Fraigniaud, T. Gonen, R. Levi, M. Medina, P. Montealegre, D. Olivetti, R. Oshman, I. Rapaport, and I. Todinca. Three Notes on Distributed Property Testing. In Proceedings 31st International Symposium on Distributed Computing (DISC), volume 91 of Leibniz International Proceedings in Informatics (LIPIcs), pages 15:1-15:30, 2017.
[14] M. Ghaffari and B. Haeupler. Distributed algorithms for planar networks I: planar embedding. In Proceedings 36th ACM Symposium on Principles of Distributed Computing (PODC), pages 29-38, 2016.
[15] M. Ghaffari and B. Haeupler. Distributed algorithms for planar networks II: low-congestion shortcuts, mst, and min-cut. In Proceedings 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 202-219, 2016.
[16] M. Ghaffari, F. Kuhn, and H.-H. Su. Distributed MST and routing in almost mixing time. In Proceedings 37th ACM Symposium on Principles of Distributed Computing (PODC), pages 131-140, 2017.
[17] M. Ghaffari and J. Li. New Distributed Algorithms in Almost Mixing Time via Transformations from Parallel Algorithms. ArXiv e-prints, May 2018.
[18] O. Goldreich and D. Ron. A sublinear bipartiteness tester for bounded degree graphs. Combinatorica, 19(3):335-373, Mar 1999.
[19] T. Gonen and R. Oshman. Lower bounds for subgraph detection in the CONGEST model. In Proceedings 21st International Conference on Principles of Distributed Systems (OPODIS), volume 95 of Leibniz International Proceedings in Informatics (LIPIcs), pages 6:1-6:16, 2018.
[20] C. Huang, D. Nanongkai, and T. Saranurak. Distributed exact weighted all-pairs shortest paths in $\tilde{O}\left(n^{5 / 4}\right)$ rounds. In Proceedings 58 th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pages 168-179, 2017.
[21] T. Izumi and F. Le Gall. Triangle finding and listing in CONGEST networks. In Proceedings 37th ACM Symposium on Principles of Distributed Computing (PODC), pages 381-389, 2017.
[22] M. Jerrum and A. Sinclair. Approximating the permanent. SIAM Journal on Computing, 18(6):1149-1178, 1989.
[23] T. Jurdziński and K. Nowicki. MST in $O(1)$ rounds of congested clique. In Proceedings 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 2620-2632, 2018.
[24] K.-i. Kawarabayashi and M. Thorup. Deterministic global minimum cut of a simple graph in near-linear time. In Proceedings 47 th Annual ACM Symposium on Theory of Computing (STOC), pages 665-674, 2015.
[25] J. H. Korhonen and J. Rybicki. Deterministic subgraph detection in broadcast CONGEST. In Proceedings 21st International Conference on Principles of Distributed Systems (OPODIS), volume 95 of Leibniz International Proceedings in Informatics (LIPIcs), pages 4:1-4:16, 2018.
[26] S. Krinninger and D. Nanongkai. A faster distributed single-source shortest paths algorithm. In Proceedings 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2018.
[27] F. Kuhn and A. R. Molla. Distributed sparse cut approximation. In Proceedings 19th International Conference on Principles of Distributed Systems (OPODIS), pages 10:1-10:14, 2015.
[28] A. Kumar, C. Seshadhri, and A. Stolman. Finding forbidden minors in sublinear time: a $O\left(n^{1 / 2+o(1)}\right)$-query one-sided tester for minor closed properties on bounded degree graphs. ArXiv e-prints, May 2018.
[29] T. C. Kwok and L. C. Lau. Finding small sparse cuts by random walk. In Proceedings 15th International Workshop on Approximation, Randomization, and Combinatorial Optimization (APPROX/RANDOM), pages 615-626, 2012.
[30] C. Lenzen. Optimal deterministic routing and sorting on the congested clique. In Proceedings 33rd ACM Symposium on Principles of Distributed Computing (PODC), pages 42-50, 2013.
[31] G. Moshkovitz and A. Shapira. Decomposing a graph into expanding subgraphs. Random Struct. Algorithms, 52(1):158-178, 2018.
[32] G. Pandurangan, P. Robinson, and M. Scquizzato. On the distributed complexity of large-scale graph computations. In Proceedings 30th ACM Symposium on Parallelism in Algorithms and Architecture (SPAA), 2018.
[33] M. Pǎtraşcu and M. Thorup. Planning for fast connectivity updates. In Proceedings 48th IEEE Symposium on Foundations of Computer Science (FOCS), pages 263-271, 2007.
[34] D. Peleg. Distributed Computing: A Locality-Sensitive Approach. SIAM, 2000.
[35] D. Peleg and V. Rubinovich. A near-tight lower bound on the time complexity of distributed minimum-weight spanning tree construction. SIAM J. Comput., 30(5):1427-1442, 2000.
[36] P. Raghavendra and D. Steurer. Graph expansion and the unique games conjecture. In Proceedings 42nd ACM Symposium on Theory of Computing (STOC), pages 755-764, 2010.
[37] A. D. Sarma, S. Holzer, L. Kor, A. Korman, D. Nanongkai, G. Pandurangan, D. Peleg, and R. Wattenhofer. Distributed verification and hardness of distributed approximation. SIAM J. Comput., 41(5):1235-1265, 2012.
[38] A. D. Sarma, A. R. Molla, and G. Pandurangan. Distributed computation of sparse cuts via random walks. In Proceedings 16th International Conference on Distributed Computing and Networking (ICDCN), pages 6:1-6:10, 2015.
[39] D. A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In Proceedings 36th Annual ACM Symposium on Theory of Computing (STOC), pages 81-90, 2004.
[40] D. A. Spielman and S.-H. Teng. A local clustering algorithm for massive graphs and its application to nearly linear time graph partitioning. SIAM J. Comput., 42(1):1-26, 2013.
[41] L. Trevisan. Approximation algorithms for unique games. Theory of Computing, 4(5):111-128, 2008.


[^0]:    ${ }^{1}$ By definition, $\operatorname{Vol}(S)=\sum_{v \in S} \operatorname{deg}(v)$.

[^1]:    ${ }^{2}$ The arboricity of a graph is the minimum number $\alpha$ such that its edge set can be partitioned into $\alpha$ forests.

[^2]:    ${ }^{3}$ Kuhn and Molla [27] further claimed that the output cut of their algorithm has balance at least $b / 2$, but this claim turns out to be incorrect (Anisur Rahaman Molla, personal communication, 2018).

[^3]:    ${ }^{4}$ There are many versions of the paper [39] available; we refer to https://arxiv.org/abs/cs/0310051v9.

[^4]:    ${ }^{5}$ Notice that when $v^{\star} \in V\left(E_{\text {out }}\right)-V_{\text {in }}, \tau_{\text {mix }}\left(G_{\text {in }}\right)=n^{o(1)}$ implies that $\tau_{\text {mix }}\left(G_{\text {in }}^{+}\right)=n^{o(1)}$ as well.
    ${ }^{6}$ In the Triangle Counting application, it is important that $v$ not count every triangle it is aware of. For example, if $v$ is assigned $\left(j, j, j^{\prime}\right), v$ knows about triangles in the subgraph induced by $V_{j}$ but should not count them.

[^5]:    ${ }^{7}$ Most of these cases do not occur in the parital orientations produced by our algorithm; nonetheless, they can occur in arbitrary partial acyclic orientations.

[^6]:    ${ }^{8}$ Of course, their lower bound holds for dense graph with $\Theta\left(n^{2}\right)$ edges; it is unclear whether $\tilde{\Omega}\left(n^{1 / 3}\right)$ is a valid lower bound for all possible $\Delta$ in the relevant range $n^{1 / 3}<\Delta<n$. The lower bound clearly does not hold when $\Delta \ll n^{1 / 3}$.

