

# Fast Distributed Algorithms for Connectivity and MST in Large Graphs

GOPAL PANDURANGAN, University of Houston  
PETER ROBINSON, Royal Holloway, University of London  
MICHELE SCQUIZZATO, University of Houston

Motivated by the increasing need to understand the algorithmic foundations of distributed large-scale graph computations, we study a number of fundamental graph problems in a message-passing model for distributed computing where  $k \geq 2$  machines jointly perform computations on graphs with  $n$  nodes (typically,  $n \gg k$ ). The input graph is assumed to be initially randomly partitioned among the  $k$  machines, a common implementation in many real-world systems. Communication is point-to-point, and the goal is to minimize the number of communication rounds of the computation.

Our main result is an (almost) optimal distributed randomized algorithm for graph connectivity. Our algorithm runs in  $\tilde{O}(n/k^2)$  rounds ( $\tilde{O}$  notation hides a polylog( $n$ ) factor and an additive polylog( $n$ ) term). This improves over the best previously known bound of  $\tilde{O}(n/k)$  [Klauck et al., SODA 2015], and is optimal (up to a polylogarithmic factor) in light of an existing lower bound of  $\tilde{\Omega}(n/k^2)$ . Our improved algorithm uses a bunch of techniques, including linear graph sketching, that prove useful in the design of efficient distributed graph algorithms. Using the connectivity algorithm as a building block, we then present fast randomized algorithms for computing minimum spanning trees, (approximate) min-cuts, and for many graph verification problems. All these algorithms take  $\tilde{O}(n/k^2)$  rounds, and are optimal up to polylogarithmic factors. We also show an almost matching lower bound of  $\tilde{\Omega}(n/k^2)$  rounds for many graph verification problems by leveraging lower bounds in random-partition communication complexity.

CCS Concepts: • **Theory of computation** → **Massively parallel algorithms; Distributed algorithms;**

Additional Key Words and Phrases: Distributed graph algorithms; graph sketching; graph connectivity; minimum spanning trees

## ACM Reference Format:

Gopal Pandurangan, Peter Robinson, and Michele Scquizzato. 2018. Fast Distributed Algorithms for Connectivity and MST in Large Graphs. *ACM Trans. Parallel Comput.* 5, 1, Article 4 (June 2018), 22 pages. <https://doi.org/10.1145/3209689>

## 1 INTRODUCTION

The focus of this paper is on distributed computation on large-scale graphs, which is becoming increasingly important with the rise of massive graphs such as the Web graph, social networks,

---

A preliminary version of this work [40] was presented at the *28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2016)*. This work was supported, in part, by US-Israel Binational Science Foundation grant 2008348, and NSF grants CCF-1527867, CCF-1540512, IIS-1633720, and CCF-1717075.

Authors' addresses: G. Pandurangan, Department of Computer Science, University of Houston, USA, [gopalpandurangan@gmail.com](mailto:gopalpandurangan@gmail.com); P. Robinson (current address), Department of Computing and Software, McMaster University, Canada, [peter.robinson@mcmaster.ca](mailto:peter.robinson@mcmaster.ca); M. Scquizzato (current address), School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Sweden, [mseq@kth.se](mailto:mseq@kth.se).

---

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

© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.  
1539-9087/2018/6-ART4 \$15.00  
<https://doi.org/10.1145/3209689>

biological networks, and other graph-structured data, and the consequent need for fast algorithms to process such graphs. The size of the graphs that we encounter in such contexts is usually so large (i.e., in the order of billions of nodes and hundreds of billions of edges) that makes their processing possible only with the use of dedicated distributed platforms. This has led to the appearance of novel frameworks for distributed computing such as MapReduce, Pregel, Giraph, and Spark. It is therefore a key task to design practical distributed algorithms with good theoretical guarantees for basic graph problems.

In this paper we study a number of fundamental graph problems in a computational model which abstracts the essence of these graph-processing systems, and present almost tight bounds on the time complexity needed to solve these problems. In this model, introduced in [24] and explained in detail in Section 1.1, the input graph is distributed across a group of  $k \geq 2$  machines that are pairwise interconnected via a communication network. The  $k$  machines jointly perform computations on an arbitrary  $n$ -vertex input graph, where typically  $n \gg k$ . The input graph is assumed to be initially randomly partitioned among the  $k$  machines, a common implementation in many real world graph processing systems (e.g., Pregel [33]). Communication is point-to-point via message passing. The computation advances in synchronous rounds, and there is a constraint on the amount of data that can cross each link of the network in each round. The goal is to minimize the time complexity, i.e., the number of rounds required by the computation. This model is aimed at investigating the amount of “speed-up” possible vis-a-vis the number of available machines, in the following sense: when  $k$  machines are used, how does the time complexity scale in  $k$ ? Which problems admit linear scaling? Is it possible to achieve super-linear scaling?

Klauck et al. [24] present lower and upper bounds for several fundamental graph problems in the  $k$ -machine model. In particular, assuming that each link has a bandwidth of one bit per round, they show a lower bound of  $\tilde{\Omega}(n/k^2)$  rounds<sup>1</sup> for the graph connectivity problem.<sup>2</sup> They also present an  $\tilde{O}(n/k)$ -round algorithm for graph connectivity and spanning tree (ST) verification. This algorithm thus exhibits a scaling linear in the number of machines  $k$ . The question of existence of a faster algorithm, and in particular of an algorithm matching the  $\tilde{\Omega}(n/k^2)$  lower bound, was left open in [24]. In this paper we answer this question affirmatively by presenting an  $\tilde{O}(n/k^2)$ -round algorithm for graph connectivity, thus achieving a speedup quadratic in  $k$ . This is optimal up to polylogarithmic (in  $n$ ) factors.

This result is important for two reasons. First, it shows that there are non-trivial graph problems for which we can obtain *superlinear* (in  $k$ ) speed-up. To elaborate further on this point, we shall take a closer look at the proof of the lower bound for the graph connectivity problem shown in [24]. Using communication complexity techniques, that proof shows that any (possibly randomized) algorithm for connectivity must exchange  $\tilde{\Omega}(n)$  bits of information across the  $k$  machines, for any  $k \geq 2$ . Since there are  $k(k-1)/2$  links in a complete network with  $k$  machines, when each link can carry  $O(\text{polylog}(n))$  bits per round, in each single round the network can deliver at most  $\tilde{O}(k^2)$  bits, and thus a lower bound of  $\tilde{\Omega}(n/k^2)$  rounds follows. The result of this paper shows that it is possible to exploit in full the available bandwidth, thus achieving a speed-up of  $\tilde{O}(k^2)$ . Second, this implies that many other important graph problems can be solved in  $\tilde{O}(n/k^2)$  rounds as well. These include computing a spanning tree, minimum spanning tree (MST), approximate min-cut, and many verification problems such as spanning connected subgraph, cycle containment, and bipartiteness.

It is important to note that under a different output requirement (explained next) there exists a  $\tilde{\Omega}(n/k)$ -round lower bound for computing a spanning tree of a graph [24], which also implies

<sup>1</sup>Throughout this paper  $\tilde{O}(f(n))$  denotes  $O(f(n) \text{polylog } n + \text{polylog } n)$ , and  $\tilde{\Omega}(f(n))$  denotes  $\Omega(f(n)/\text{polylog } n)$ .

<sup>2</sup>If each link has bandwidth of  $B$  bits per round, then this bound generalizes to  $\tilde{\Omega}(n/Bk^2)$  rounds.

the same lower bound for other fundamental problems such as computing an MST, breadth-first tree, and shortest paths tree. However, this lower bound holds under the requirement that each vertex (i.e., the machine which hosts the vertex) must know at the end of the computation the “status” of all of its incident edges, that is, whether they belong to an ST or not, and output their respective status. (This is the output criterion usually required for distributed algorithms [32, 43].) The proof of the lower bound exploits this criterion to show that any algorithm requires some machine receiving  $\Omega(n)$  bits of information, and since any machine has  $k - 1$  incident links, this results in a  $\tilde{\Omega}(n/k)$  lower bound. On the other hand, if we relax the output criterion to require the final status of each edge to be known by *some* machine, then we show that this can be accomplished in  $\tilde{O}(n/k^2)$  rounds using the fast connectivity algorithm of this paper.

### 1.1 The Model

We now describe the adopted model of distributed computation, the *k-machine model* (a.k.a. the *Big Data model*), introduced in [24] and further investigated in [5, 10, 42, 45]. The model consists of a set of  $k \geq 2$  machines  $N = \{M_1, M_2, \dots, M_k\}$  that are pairwise interconnected by bidirectional point-to-point communication links. Each machine executes an instance of a distributed algorithm. The computation advances in synchronous rounds where, in each round, machines can exchange messages over their communication links and perform some local computation. Machines do not share any memory and have no other means of communication. Each link is assumed to have a bandwidth of  $O(\text{polylog}(n))$  bits per round, where  $n$  is the input size, i.e.,  $O(\text{polylog}(n))$  bits can be transmitted over each link in each round. (As discussed in [24, Theorem 4.1], it is easy to rewrite bounds to scale in terms of the actual inter-machine bandwidth.) There is an alternate (but equivalent) way to view this communication restriction: instead of putting a bandwidth restriction on the links, we can put a restriction on the amount of information that each *machine* can communicate (i.e., send/receive) in each round. The results we obtain in the bandwidth-restricted model also apply to the latter model [24]. Local computation within a machine is considered to happen instantaneously at zero cost, while the exchange of messages between machines is the costly operation. (However, we note that in all the algorithms of this paper, every machine in every round performs a computation bounded by a polynomial in  $n$ .) We assume that each machine has access to a private source of true random bits.

Although the *k-machine model* is a general model of distributed computation that can be applied to study any (large-scale data) problem, in this paper we aim at investigating graph problems in it. Specifically, we are given an input graph  $G$  with  $n$  vertices, each associated with a unique integer ID from  $[n]$ , and  $m$  edges. To avoid trivialities, we will assume that  $n \geq k$  (typically,  $n \gg k$ ). Initially, the entire graph  $G$  is not known by any single machine, but rather partitioned among the  $k$  machines in a “balanced” fashion, i.e., the nodes and/or edges of  $G$  are partitioned approximately evenly among the machines. We assume a *vertex-partition* model, whereby vertices, along with information of their incident edges, are partitioned across machines. Specifically, the type of partition that we will assume throughout is the *random vertex partition (RVP)*, that is, each vertex of the input graph is assigned randomly to one machine.<sup>3</sup> (This is the typical way used by many real systems, such as Pregel [33], to initially distribute the input graph among the machines. See also [9, 46].) However, we notice that our upper bounds also hold under the much weaker assumption whereby it is only required that nodes and edges of the input graph are partitioned approximately evenly among the

<sup>3</sup>In Section 1.3 we will discuss an alternate partitioning model, the *random edge partition (REP)* model, where each edge of  $G$  is assigned independently and randomly to one of the  $k$  machines, and show how the results in the random vertex partition model can be related to the random edge partition model.

machines; on the other hand, lower bounds under RVP clearly apply to worst-case partitions as well.

More formally, in the random vertex partition variant, each vertex of  $G$  is assigned independently and uniformly at random to one of the  $k$  machines. If a vertex  $v$  is assigned to machine  $M_i$  we say that  $M_i$  is the *home machine* of  $v$  and, with a slight abuse of notation, write  $v \in M_i$ . When a vertex is assigned to a machine, all its incident edges are assigned to that machine as well; i.e., the home machine knows the IDs of the neighbors of that vertex as well as the identity of the home machines of the neighboring vertices (and the weights of the corresponding edges in case  $G$  is weighted). Note that an immediate property of the RVP model is that the number of vertices at each machine is *balanced*, i.e., each machine is the home machine of  $\tilde{\Theta}(n/k)$  vertices with high probability. A convenient way to implement the RVP model is through hashing: each vertex (ID) is hashed to one of the  $k$  machines. Hence, if a machine knows a vertex ID, it also knows where it is hashed to.

Eventually, each machine  $M_i$  must set a designated local output variable  $o_i$  (which need not depend on the set of vertices assigned to  $M_i$ ), and the *output configuration*  $o = \langle o_1, \dots, o_k \rangle$  must satisfy the feasibility conditions of the problem at hand. For example, for the minimum spanning tree problem each  $o_i$  corresponds to a set of edges, and the edges in the union of such sets must form an MST of the input graph.

In this paper, we show results for distributed Monte Carlo algorithms. Recall that a Monte Carlo algorithm is a randomized algorithm whose output may be incorrect with some probability. Formally, we say that an algorithm computes a function  $f$  with  $\epsilon$ -error if for every input it outputs the correct answer with probability at least  $1 - \epsilon$ , where the probability is over the random partition and the random bit strings used by the algorithm (if any). The *round (time) complexity* of an algorithm is the maximum number of communication rounds until termination. For any  $n$  and problem  $\mathcal{P}$  on  $n$ -node graphs, we let the *time complexity of solving  $\mathcal{P}$  with  $\epsilon$  error probability* in the  $k$ -machine model, denoted by  $\mathcal{T}_\epsilon(\mathcal{P})$ , be the minimum  $T(n)$  such that there exists an  $\epsilon$ -error protocol that solves  $\mathcal{P}$  and terminates in  $T(n)$  rounds. For any  $0 \leq \epsilon \leq 1$ , graph problem  $\mathcal{P}$  and function  $T : \mathbb{Z}_+ \rightarrow \mathbb{Z}_+$ , we say that  $\mathcal{T}_\epsilon(\mathcal{P}) = O(T(n))$  if there exists an integer  $n_0$  and a real  $c$  such that for all  $n \geq n_0$ ,  $\mathcal{T}_\epsilon(\mathcal{P}) \leq cT(n)$ . Similarly, we say that  $\mathcal{T}_\epsilon(\mathcal{P}) = \Omega(T(n))$  if there exists an integer  $n_0$  and a real  $c$  such that for all  $n \geq n_0$ ,  $\mathcal{T}_\epsilon(\mathcal{P}) \geq cT(n)$ . For our upper bounds, we will usually have  $\epsilon = 1/n$ , which means high probability algorithms, i.e., succeeding with probability at least  $1 - 1/n$ . In this case, we will sometimes omit  $\epsilon$  and simply say the time bound applies “with high probability.”

## 1.2 Contributions and Techniques

The main result of this paper, presented in Section 2, is a randomized Monte Carlo algorithm for the  $k$ -machine model that determines the connected components of an undirected graph  $G$  correctly with high probability and that terminates in  $\tilde{O}(n/k^2)$  rounds with high probability.<sup>4</sup> This improves upon the previous best bound of  $\tilde{O}(n/k)$  [24], since it is strictly superior in the wide range of parameter  $k = \Theta(n^\epsilon)$ , for all constants  $\epsilon \in (0, 1)$ . Improving over this bound is non-trivial since various attempts to design a faster connectivity algorithm fail due to the fact that they end up creating congestion at a particular machine, i.e., up to  $n$  bits may need to be sent/received by a machine, leading to a  $\tilde{O}(n/k)$  bound at best. For example, a simple algorithm for connectivity is simply flooding: each vertex floods the lowest labeled vertex that it has seen so far; at the end each vertex will have the label of the lowest labeled vertex in its component.<sup>5</sup> It can be shown that

<sup>4</sup>Since the focus is on the scaling of the time complexity with respect to  $k$ , we omit the polylogarithmic factors in our time bounds. Moreover, the cumulative polylogarithmic factor is not large:  $O(\log^3 n)$ .

<sup>5</sup>This algorithm has been implemented in a variant of Giraph [48].

the above algorithm takes  $\Theta(n/k + D)$  rounds (where  $D$  is the graph diameter) in the  $k$ -machine model by using the Conversion Theorem of [24]. Hence new techniques are needed to break the  $n/k$ -round barrier.

Our connectivity algorithm is the result of the application of the following three techniques.

1. *Randomized Proxy Computation*. This technique, similar to known techniques used in randomized routing algorithms [49], is used to load-balance congestion at any given machine by redistributing it evenly across the  $k$  machines. This is achieved, roughly speaking, by re-assigning the executions of individual nodes uniformly at random among the machines. It is crucial to distribute *both* computation and communication across the machines to avoid congestion at any particular machine. In fact, this allows one to move away from the communication pattern imposed by the topology of the input graph (which can cause congestion at a particular machine) to a more balanced communication.

2. *Distributed Random Ranking (DRR)*. DRR [8] is a simple technique that will be used to build trees of low height in the connectivity algorithm. Our connectivity algorithm is divided into phases, in each of which we do the following: each current component (in the first phase, each vertex is a component by itself) chooses one *outgoing edge* and then components are combined by merging them along outgoing edges. If done naively, this may result in a long chain of merges, resulting in a component tree of high diameter; thus, communication along this tree would take a long time. To avoid this we resort to DRR, which suitably reduces the number of merges. With DRR, each component chooses a random *rank*, which is simply a random number, say in the interval  $[1, n^3]$ ; a component  $C_i$  then merges with the component  $C_j$  on the other side of its selected outgoing edge if and only if the rank of  $C_j$  is larger than the rank of  $C_i$ . Otherwise,  $C_i$  does not merge with  $C_j$ , and thus it becomes the root of a DRR tree, which is a tree induced by the components and the set of the outgoing edges that have been used in the above merging procedure. It can be shown that the height of a DRR tree is bounded by  $O(\log n)$  with high probability.

3. *Linear Graph Sketching*. Linear graph sketching [1, 2, 34] is crucially helpful in efficiently finding an outgoing edge of a component. A *sketch* for a vertex (or a component) is a short (say,  $O(\text{polylog } n)$  bits) bit vector that efficiently encodes the adjacency list of the vertex. Sampling from this sketch gives a random (outgoing) edge of this vertex (component). A very useful property is the linearity of the sketches: adding the sketches of a set of vertices gives the sketch of the component obtained by combining the vertices; the edges between the vertices (i.e., the intra-component edges) are automatically “cancelled”, leaving only a sketch of the outgoing edges. Linear graph sketches were originally used to process dynamic graphs in the (centralized) streaming model [1, 2, 34]. Here, in a distributed setting, we use them to reduce the amount of communication needed to find an outgoing edge; in particular, graph sketches will avoid us from checking whether an edge is an inter-component or an intra-component edge, and this will crucially reduce communication across machines. We note that earlier distributed algorithms such as the classical GHS algorithm [15] for the MST problem would incur too much communication since they involve checking the status of each edge of the graph.

We observe that it does not seem straightforward to effectively exploit these techniques in the  $k$ -machine model: for example, linear sketches can be easily applied in the distributed streaming model by sending to a coordinator machine the sketches of the partial stream, which then will be added to obtain the sketch of the entire stream. Mimicking this trivial strategy in the  $k$ -machine model would cause too much congestion at one node, leading to a  $\tilde{O}(n/k)$  time bound.

Using the above techniques and the fast connectivity algorithm, in Section 3 we give algorithms for many other important graph problems. In particular, we present a  $\tilde{O}(n/k^2)$ -round algorithm for computing an MST (and hence an ST). We also present  $\tilde{O}(n/k^2)$ -round algorithms for approximate

min-cut, and for many graph verification problems including spanning connected subgraph, cycle containment, and bipartiteness. All these algorithms are optimal up to a polylogarithmic factor.

In Section 4 we show a lower bound of  $\tilde{\Omega}(n/k^2)$  rounds for many verification problems by simulating the  $k$ -machine model in a 2-party model of communication complexity where the inputs are randomly assigned to the players.

### 1.3 Related Work

The theoretical study of large-scale graph computations in distributed systems is relatively new. Several works have been devoted to developing algorithm for MapReduce (see, e.g., [22, 27, 30] and references therein). We note that the flavor of the theory developed for MapReduce is quite different compared to the one for the  $k$ -machine model. Minimizing communication is also the key goal in MapReduce algorithms; however this is usually achieved by making sure that the data is made small enough quickly (that is, in a small number of MapReduce rounds) to fit into the memory of a single machine (see, e.g., the MapReduce algorithm for MST in [27]).

For a comparison of the  $k$ -machine model with other models for parallel and distributed processing, including Bulk-Synchronous Parallel (BSP) model [50], MapReduce [22], and the congested clique, we refer to [51]. In particular, according to [51], “Among all models with restricted communication the “big data” [ $k$ -machine] model is the one most similar to the MapReduce model”.

The  $k$ -machine model is closely related to the BSP model; it can be considered as a simplified version of BSP, where the costs of local computation and of synchronization (which happens at the end of every round) are ignored. Unlike the BSP and refinements thereof, which have several different parameters that make the analysis of algorithms complicated [51], the  $k$ -machine model is characterized by just one parameter, the number of machines; this makes the model simple enough to be analytically tractable, thus easing the job of designing and analyzing algorithms, while at the same time it still captures the key features of large-scale distributed computations.

The  $k$ -machine model is also related to the classical CONGEST model [43], and in particular to the *congested clique* model, which recently has received considerable attention (see, e.g., [7, 13, 17, 18, 28, 29, 31, 36]). The main difference is that the  $k$ -machine model is aimed at the study of large-scale computations, where the size  $n$  of the input is significantly bigger than the number of available machines  $k$ , and thus many vertices of the input graph are mapped to the same machine, whereas the two aforementioned models are aimed at the study of distributed network algorithms, where  $n = k$  and each vertex corresponds to a dedicated machine. More “local knowledge” is available per vertex (since it can access for free information about other vertices in the same machine) in the  $k$ -machine model compared to the other two models. On the other hand, all vertices assigned to a machine have to communicate through the links incident on this machine, which can limit the bandwidth (unlike the other two models where each vertex has a dedicated processor). These differences manifest in the time complexity. In particular, the fastest known distributed algorithm in the congested clique model for a given problem may not give rise to the fastest algorithm in the  $k$ -machine model. For example, the fastest algorithms for MST in the congested clique model ([17, 18, 31]) require  $\Theta(n^2)$  messages; implementing these algorithms in the  $k$ -machine model requires  $\Theta(n^2/k^2)$  rounds. Conversely, the slower GHS algorithm [15] gives an  $\tilde{O}(n/k)$  bound in the  $k$ -machine model. The recently developed techniques (see, e.g., [12–14, 18, 37]) used to prove time lower bounds in the CONGEST model and in the congested clique model are not directly applicable here.

The work closest in spirit to ours is the recent work of Woodruff and Zhang [52]. This paper considers a number of basic statistical and graph problems in a distributed message-passing model similar to the  $k$ -machine model. However, there are some important differences. First, their model

is asynchronous, and the cost function is the communication complexity, which refers to the total number of bits exchanged by the machines during the computation. Second, a *worst-case* distribution of the input is assumed, while we assume a random distribution. Third, which is an important difference, they assume an edge partition model for the problems on graphs, that is, the edges of the graph (as opposed to its vertices) are partitioned across the  $k$  machines. In particular, for the connectivity problem, they show a message complexity lower bound of  $\tilde{\Omega}(nk)$  which essentially translates to a  $\tilde{\Omega}(n/k)$  round lower bound in the  $k$ -machine model (assuming an initial worst-case edge partition); it can be shown by using their proof technique that this lower bound also applies to the *random edge partition (REP)* model, where edges are partitioned randomly among machines, as well. On the other hand, it is easy to show an  $\tilde{O}(n/k)$  upper bound for connectivity and MST in the REP model.<sup>6</sup> Hence, in the REP model,  $\tilde{\Theta}(n/k)$  is a tight bound for connectivity and other related problems such as MST. In contrast, in the RVP model (arguably, a more natural partition model), we show that  $\tilde{\Theta}(n/k^2)$  is the tight bound.

From the technical point of view, [23] also uses an idea similar to linear sketching. Their technique might also be useful in the context of the  $k$ -machine model.

## 2 THE CONNECTIVITY ALGORITHM

In this section we present our main result, a Monte Carlo randomized algorithm for the  $k$ -machine model that determines the connected components of an undirected graph  $G$  correctly with high probability and that terminates in  $\tilde{O}(n/k^2)$  rounds with high probability. This algorithm is optimal, up to polylog( $n$ )-factors, by virtue of a lower bound of  $\tilde{\Omega}(n/k^2)$  rounds [24].

Before delving into the details of our algorithm, as a warm-up we briefly discuss simpler, but less efficient, approaches. The easiest way to solve *any* problem in our model is to first collect all available graph data at a single machine and then solve the problem locally. For example, one could first elect a referee among the machines, which requires  $O(1)$  rounds [26], and then instruct every machine to send its local data to the referee machine. Since the referee machine needs to receive  $O(m)$  information in total but has only  $k - 1$  links of bounded bandwidth, this requires  $\Omega(m/k)$  rounds.

A more refined approach to obtain a distributed algorithm for the  $k$ -machine model is to use the Conversion Theorem of [24], which provides a simulation of a congested clique algorithm  $\mathcal{A}$  in  $\tilde{O}(M/k^2 + \Delta'T/k)$  rounds in the  $k$ -machine model, where  $M$  is the message complexity of  $\mathcal{A}$ ,  $T$  is its round complexity, and  $\Delta'$  is an upper bound to the total number of messages sent (or received) by a single node in a single round. (All these parameters refer to the performance of  $\mathcal{A}$  in the congested clique model.) Unfortunately, existing algorithms (e.g., [15, 47]) either require  $\Delta'$  to scale to the maximum node degree, and thus the converted time complexity bound in the  $k$ -machine model is no better than  $\tilde{O}(n/k)$ , or, as is the case of the recent fast algorithms [17, 18, 31], have a high ( $\Theta(n^2)$ ) message complexity and thus cannot run in less than  $\Theta(n^2/k^2)$  rounds in the  $k$ -machine model. The recent MST algorithm of Pemmaraju and Sardeshmukh [44] is both message- and time-efficient in the congested clique: it uses  $\tilde{O}(n)$  messages and runs in  $O(\log n)$  rounds. However, there are several steps in this algorithm where a single node receives  $\Theta(n)$  messages, and thus directly applying the Conversion Theorem of [24] to this algorithm gives only an  $\tilde{O}(n/k)$ -round algorithm in the  $k$ -machine model. Therefore, in order to break the  $\tilde{\Omega}(n/k)$  barrier, we must develop new techniques that directly exploit the additional locality available in the  $k$ -machine model.

<sup>6</sup>The high-level idea of the MST algorithm in the REP model is: (1) First “filter” the edges assigned to one machine using the cut and cycle properties of a MST [21]; this leaves each machine with  $O(n)$  edges; (2) Convert this edge distribution to a RVP which can be accomplished in  $\tilde{O}(n/k)$  rounds via hashing the vertices randomly to machines and then routing the edges appropriately; then apply the RVP bound.

In the next subsection we give a high level overview of our algorithm, and then formally present all the technical details in the subsequent subsections.

## 2.1 Overview of the Algorithm

Our algorithm follows a Borůvka-style strategy [6], that is, it repeatedly merges adjacent *components* of the input graph  $G$ , which are connected subgraphs of  $G$ , to form larger (connected) components. The output of each of these phases is a *labeling* of the nodes of  $G$  such that nodes that belong to the same current component have the same label. At the beginning of the first phase, each node is labeled with its own unique ID, forms a distinct component, and is also the *component proxy* of its own component. Note that, at any phase, a component contains up to  $n$  nodes, which might be spread across different machines; we use the term *component part* to refer to all those nodes of the component that are held by the same machine. Hence, at any phase every component is partitioned in at most  $k$  component parts. At the end of the algorithm each vertex has a label such that any two vertices have the same label if and only if they belong to the same connected component of  $G$ .

Our algorithm relies on *linear graph sketches* as a tool to enable communication-efficient merging of multiple components. Intuitively speaking, a (random) linear sketch  $\mathbf{s}_u$  of a node  $u$ 's graph neighborhood returns a sample chosen uniformly at random from  $u$ 's incident edges. Interestingly, such a linear sketch can be represented as matrices using only  $O(\text{polylog}(n))$  bits [19, 34]. A crucial property of these sketches is that they are linear: that is, given sketches  $\mathbf{s}_u$  and  $\mathbf{s}_v$ , the *combined sketch*  $\mathbf{s}_u + \mathbf{s}_v$  (“+” refers to matrix addition) has the property that, w.h.p., it yields a random sample of the edges incident to  $(u, v)$  in a graph where we have contracted the edge  $(u, v)$  to a single node. We describe the technical details in Section 2.3.

We now describe how to communicate these graph sketches in an efficient manner: Consider a component  $C$  that is split into  $j$  parts  $P_1, P_2, \dots, P_j$ , the nodes of which are hosted at machines  $M_1, M_2, \dots, M_j$ . To find an outgoing edge for  $C$ , we first instruct each machine  $M_i$  to construct a linear sketch of the graph neighborhood of each of the nodes in part  $P_i$ . Then, we sum up these  $|P_i|$  sketches, yielding a sketch  $\mathbf{s}_{P_i}$  for the neighborhood of part  $P_i$ . To combine the sketches of the  $j$  distinct parts, we now select a *random component proxy* machine  $M_{C,r}$  for the current component  $C$  at round  $r$  (see Section 2.2). Next, machine  $M_i$  sends  $\mathbf{s}_{P_i}$  to machine  $M_{C,r}$ ; note that this causes at most  $k$  messages to be sent to the component proxy. Finally, machine  $M_{C,r}$  computes  $\mathbf{s}_C = \sum_{i=1}^j \mathbf{s}_{P_i}$ , and then uses  $\mathbf{s}_C$  to sample an edge incident to some node in  $C$ , which, by construction, is guaranteed to have its endpoint in a distinct component  $C'$ . (See Section 2.4.)

At this point, each component proxy has sampled an inter-component edge inducing the edges of a *component graph*  $C$  where each vertex corresponds to a component. To enable the efficient merging of components, we employ the distributed random ranking technique of [8] to break up any long paths of  $C$  into more manageable directed trees of depth  $O(\log n)$ . To this end, each component chooses a rank independently and uniformly at random from  $[0, 1]$ , and then (virtually) connects to its neighboring component (according to  $C$ ) via a (conceptual) directed edge if and only if the latter has a higher rank. Thus, this process results in a collection of disjoint rooted trees, rooted at the node of highest (local) rank.

The merging of the components of each tree  $\mathcal{T}$  proceeds from the leaves upward (in parallel for each tree). In the first merging phase, each leaf  $C_j$  of  $\mathcal{T}$  merges with its parent  $C'$  by relabeling the component labels of all of their nodes with the label of  $C'$ . Note that the proxy  $M_{C_j}$  knows the labeling of  $C'$ , as it has computed the outgoing edge from a vertex in  $C_j$  to a vertex in  $C'$ . Therefore, machine  $M_{C_j}$  sends the label of  $C_j$  to all the machines that hold a part of  $C_j$ . In Section 2.5 we show that this can be done in parallel (for all leaves of all trees) in  $\tilde{O}(n/k^2)$  rounds.

Finally, in Section 2.6 we prove that  $O(\log n)$  repetitions of the above process suffice to ensure that the components at the end of the last phase correspond to the connected components of the input graph  $G$ .

## 2.2 Communication via Random Proxy Machines

Recall that our algorithm iteratively groups vertices into components and subsequently merges such components according to the topology of  $G$ . Each of these components may be split into multiple component parts spanning multiple machines. Hence, to ensure efficient load balancing of the messages that machines need to send on behalf of the component parts that they hold, the algorithm performs all communication via *proxy machines*.

Our algorithm proceeds in phases, and each phase consists of iterations. Consider the  $\rho$ -th iteration of the  $j$ -th phase of the algorithm, with  $\rho, j \geq 1$ . We construct a “sufficiently” random hash function  $h_{j,\rho}$ , such that, for each component  $C$ , the machine with ID  $h_{j,\rho}(C) \in [k]$  is selected as the proxy machine for component  $C$ . First, machine  $M_1$  generates  $\ell = \tilde{\Theta}(n/k)$  random bits from its private source of randomness.  $M_1$  will distribute these random bits to all other machines via the following simple routing mechanism that proceeds in sequences of two rounds.  $M_1$  selects  $k$  bits  $b_1, b_2, \dots, b_{k-1}$  from the set of its  $\ell$  private random bits that remain to be distributed, and sends bit  $b_i$  across its  $i$ -th link to machine  $M_{i+1}$ . Upon receiving  $b_i$ , machine  $M_{i+1}$  broadcasts  $b_i$  to all machines in the next round. This ensures that bits  $b_1, b_2, \dots, b_{k-1}$  become common knowledge within two rounds. Repeating this process to distribute all the  $\ell = \tilde{\Theta}(n/k)$  bits takes  $\tilde{O}(n/k^2)$  rounds, after those all the machines have the  $\ell$  random bits generated by  $M_1$ . We leverage a result of [3] (cf. in its formulation as Theorem 2.1 in [4]), which tells us that we can generate a random hash function such that it is  $d$ -wise independent by using only  $O(d \log n)$  true random bits. We instruct machine  $M_1$  to disseminate  $d = \ell \log n = n \text{ polylog}(n)/k$  of its random bits according to the above routing process and then each machine locally constructs the same hash function  $h_{j,\rho}$ , which is then used to determine the component proxies throughout iteration  $\rho$  of phase  $j$ .

We now show that communication via such proxy machines is fast in the  $k$ -machine model.

**LEMMA 2.1.** *Suppose that each machine  $M$  generates a message of size  $O(\text{polylog}(n))$  bits for each component part residing on  $M$ ; let  $m_i$  denote the message of part  $P_i$  and let  $C$  be the component of which  $P_i$  is a part. If each  $m_i$  is addressed to the proxy machine  $M_C$  of component  $C$ , then all messages are delivered within  $\tilde{O}(n/k^2)$  rounds with high probability.*

**PROOF.** Observe that, except for the very first phase of the algorithm, the claim does not immediately follow from a standard balls-into-bins argument because not all the destinations of the messages are chosen independently and uniformly at random, as any two distinct messages of the same component have the same destination.

Let us stipulate that any component part held by machine  $M_i$  is the  $i$ -th component part of its component, and denote this part with  $P_{i,j}$ ,  $i \in [k]$ ,  $j \in [n]$ , where  $P_{i,j} = \emptyset$  means that in machine  $i$  there is no component part for component  $j$ . Suppose that the algorithm is in phase  $j'$  and iteration  $\rho$ . By construction, the hash function  $h_{j',\rho}$  is  $\tilde{\Theta}(n/k)$ -wise independent, and all the component parts held by a single machine are parts of different components. Since  $M_i$  has at most  $\tilde{\Theta}(n/k)$  distinct component parts w.h.p., it follows that all the proxy machines selected by the component parts held by machine  $M_i$  are distributed independently and uniformly at random. Let  $y$  be the number of distinct component parts held by a machine  $M_i$  that is,  $y = |\{P_{i,j} : P_{i,j} \neq \emptyset\}| = \tilde{O}(n/k)$  (w.h.p.).

Consider a link of  $M_i$  connecting it to another machine  $M_1$ . Let  $X_t$  be the indicator variable that takes value 1 if  $M_1$  is the component proxy of part  $t$  (of  $M_i$ ), and let  $X_t = 0$  otherwise. Let  $X = \sum_{i=1}^y X_i$  be the number of component parts that chose their proxy machine at the endpoint of

link  $(M_i, M_1)$ . Since  $\Pr(X_i = 1) = 1/(k-1)$ , we have that the expected number of messages that have to be sent by this machine over any specific link is  $E[X] = y/(k-1)$ .

First, consider the case  $y \geq 11k \log n$ . As the  $X_i$ 's are  $\tilde{\Theta}(n/k)$ -wise independent, all proxies by the component parts of  $M_i$  are chosen independently and thus we can apply a standard Chernoff bound (see, e.g., [35]), which gives

$$\Pr\left(X \geq \frac{7y}{4(k-1)}\right) \leq e^{-3y/16(k-1)} < e^{-\frac{2k \log n}{k}} < \frac{1}{n^2}.$$

By applying the union bound over the  $k \leq n$  machines we conclude that w.h.p. every machine sends  $\tilde{O}(n/k^2)$  messages to each proxy machine, and this requires  $\tilde{O}(n/k^2)$  rounds.

Consider now the case  $y < 11k \log n$ . It holds that  $6E[X] = 6y/(k-1) < 6 \cdot 11k \log n/(k-1) \leq 132 \log n$ , and thus, by a standard Chernoff bound [35, Theorem 4.4 (3)],

$$\Pr(X \geq 132 \log n) \leq 2^{-132 \log n} = \frac{1}{n^{132}}.$$

Applying the union bound over the  $k \leq n$  machines yields the result.  $\square$

### 2.3 Linear Graph Sketches

As we will see in Section 2.5, our algorithm proceeds by merging components across randomly chosen inter-component edges. In this subsection we show how to provide these sampling capabilities in a communication-efficient way in the  $k$ -machine model by implementing random linear graph sketches. Our description follows the notation of [34].

Recall that each vertex  $u$  of  $G$  is associated with a unique integer ID from  $[n]$  (known to its home machine) which, for simplicity, we also denote by  $u$ .<sup>7</sup> For each vertex  $u$  we define the *incidence vector*  $\mathbf{a}_u \in \{-1, 0, 1\}^{\binom{n}{2}}$  of  $u$ , which describes the incident edges of  $u$ , as follows:

$$\mathbf{a}_u[(x, y)] = \begin{cases} 1 & \text{if } u = x < y \text{ and } (x, y) \in E, \\ -1 & \text{if } x < y = u \text{ and } (x, y) \in E, \\ 0 & \text{otherwise.} \end{cases}$$

Note that the vector  $\mathbf{a}_u + \mathbf{a}_v$  corresponds to the incidence vector of the contracted edge  $(u, v)$ . Intuitively speaking, summing up incidence vectors “zeroes out” edges between the corresponding vertices, hence the vector  $\sum_{u \in C} \mathbf{a}_u$  represents the outgoing edges of a component  $C$ .

Since each incidence vector  $\mathbf{a}_u$  requires polynomial space, it would be inefficient to directly communicate vectors to component proxies. Instead, we construct a random linear sketch  $\mathbf{s}_u$  of  $\text{polylog}(n)$ -size that has the property of allowing us to sample uniformly at random a nonzero entry of  $\mathbf{a}_u$  (i.e., an edge incident to  $u$ ). (This is referred to as  $\ell_0$ -sampling in the streaming literature, see e.g. [34].) It is shown in [19] that  $\ell_0$ -sampling can be performed by linear projections. Therefore, at the beginning of each phase  $j$  of our algorithm, we instruct each machine to create a new (common)  $\text{polylog}(n) \times \binom{n}{2}$  sketch matrix  $L_j$ , which we call *phase  $j$  sketch matrix*.<sup>8</sup> Then, each machine  $M$  creates a sketch  $\mathbf{s}_u = L_j \cdot \mathbf{a}_u$  for each vertex  $u$  that resides on  $M$ . Hence, each  $\mathbf{s}_u$  can be represented by a polylogarithmic number of bits.

Observe that, by linearity, we have  $L_j \cdot \mathbf{a}_u + L_j \cdot \mathbf{a}_v = L_j \cdot (\mathbf{a}_u + \mathbf{a}_v)$ . In other words, a crucial property of sketches is that the sum  $\mathbf{s}_u + \mathbf{s}_v$  is itself a sketch that allows us to sample an edge incident to the contracted edge  $(u, v)$ . We summarize these properties in the following statement.

<sup>7</sup>Note that the asymptotics of our results do not change if the size of the ID space is  $O(\text{polylog}(n))$ .

<sup>8</sup>Here we describe the construction as if nodes have access to a source of shared randomness (to create the sketch matrix). In Section 2.3.1 we show how to remove this assumption.

LEMMA 2.2. Consider a phase  $j$ , and let  $P$  a subgraph of  $G$  induced by vertices  $\{u_1, \dots, u_\ell\}$ . Let  $s_{u_1}, \dots, s_{u_\ell}$  be the associated sketches of vertices in  $P$  constructed by applying the phase  $j$  sketch matrix to the respective incidence vectors. Then, the combined sketch  $s_P = \sum_{i=1}^{\ell} s_{u_i}$  can be represented using  $O(\text{polylog}(n))$  bits and, by querying  $s_P$ , it is possible (w.h.p.) to sample a random edge incident to  $P$  (in  $G$ ) that has its other endpoint in  $G \setminus P$ .

2.3.1 *Constructing Linear Sketches Without Shared Randomness.* Our construction of the linear sketches described so far requires  $\tilde{O}(n)$  fully independent random bits that would need to be shared by all machines. However, it is shown in [11, Theorem 1] that there exists an  $\ell_0$ -sampler, having the same linearity properties, that uses  $\Theta(n)$  random bits which are only  $\Theta(\log n)$ -wise independent. Analogously as in Section 2.2, we can generate the required  $\Theta(\log^2 n)$  true random bits at machine  $M_1$ , distribute them among all other machines in  $O(1)$  rounds, and then invoke Theorem 2.1 of [4] at each machine in parallel to generate the required (shared)  $\Theta(\log n)$ -wise independent random bits for constructing the sketches.

## 2.4 Outgoing Edge Selection

Now that we know how to construct a sketch of the graph neighborhood of any set of vertices, we will describe how to combine these sketches in a communication-efficient way in the  $k$ -machine model. The goal of this step is, for each (current) component  $C$ , to find an outgoing edge that connects  $C$  to some other component  $C'$ .

Recall that  $C$  itself might be split into parts  $P_1, P_2, \dots, P_j$  across multiple machines. Therefore, as a first step, each machine  $M_i$  locally constructs the combined sketch for each part that resides in  $M_i$ . By Lemma 2.2, the resulting sketches have polylogarithmic size each and present a sketch of the incidences of their respective component parts. Next, we combine the sketches of the individual parts of each component  $C$  to a sketch of  $C$ , by instructing the machines to send the sketch of each part  $P_i$  (of component  $C$ ) to the proxy machine of  $C$ . By Lemma 2.1, all of these messages are delivered to the component proxies within  $\tilde{O}(n/k^2)$  rounds. Finally, the component proxy machine of  $C$  combines the received sketches to yield a sketch of  $C$ , and randomly samples an outgoing edge of  $C$  (see Lemma 2.2). Thus, at the end of this procedure, every component (randomly) selected exactly one neighboring component. We now show that the complexity of this procedure is  $\tilde{O}(n/k^2)$  w.h.p.

LEMMA 2.3. Every component can select exactly one outgoing edge in  $\tilde{O}(n/k^2)$  rounds with high probability.

PROOF. Clearly, since at every moment each node has a unique component's label, each machine holds  $\tilde{O}(n/k)$  component's parts w.h.p. Each of these parts selected at most one edge, and thus each machine "selected"  $\tilde{O}(n/k)$  edges w.h.p. All these edges have to be sent to the corresponding proxy. By Lemma 2.1, this requires  $\tilde{O}(n/k^2)$  rounds.

The procedure is completed when the proxies communicate the decision to each of the at most  $k$  components' parts. This entails as many messages as in the first part to be routed using exactly the same machines' links used in the first part, with the only difference being that messages now travel in the opposite direction. The lemma follows.  $\square$

## 2.5 Merging of Components

After the proxy machine of each component  $C$  has selected one edge connecting  $C$  to a different component, all the neighboring components are merged so as to become a new, bigger component. This is accomplished by relabeling the nodes of the graph such that all the nodes in the same (new) component have the same label. Notice that the merging is thus only virtual, that is, component

parts that compose a new component are not moved to a common machine; rather, nodes (and their incident edges) remain in their home machine, and just get (possibly) assigned a new label.

We can think of the components along with the sampled outgoing edges as a *component graph*  $C$ . We use the *distributed random ranking* (DRR) technique [8] to avoid having long chains of components (i.e., long paths in  $C$ ). That is, we will (conceptually) construct a forest of directed trees that is a subgraph (modulo edge directions) of the component graph  $C$  and where each tree has depth  $O(\log n)$ .<sup>9</sup> The component proxy of each component  $C$  chooses a rank independently and uniformly at random from  $[0, 1]$ . (It is easy to see that an accuracy of  $\Theta(\log n)$  bits suffices to break ties w.h.p.) Then, the proxy machine of  $C$  (virtually) connects  $C$  to its neighboring component  $C'$  if and only if the rank chosen by the latter's proxy is higher. In this case, we say that  $C'$  becomes the *parent* of  $C$  and  $C$  is a *child* of  $C'$ .

LEMMA 2.4. *After  $\tilde{O}(n/k^2)$  rounds, the structure of the DRR-tree is completed with high probability.*

PROOF. We need to show that every proxy machine of a non-root component knows its higher-rank parent component and that every root proxy machine knows that it is root. Note that during this step the proxy machines of the child components communicate with the respective parent proxy machines. Moreover, the number of messages sent for determining the ordering of the DRR-trees is guaranteed to be  $O(n)$  with high probability, since  $C$  has only  $O(n)$  edges. By Lemma 2.1, it follows that the delivery of these messages can be completed in  $\tilde{O}(n/k^2)$  rounds w.h.p.

Since links are bidirectional, the parent proxies are able to send their replies within the same number of rounds, by re-running the message schedule of the child-to-parent communication in reverse order.  $\square$

If a component has the highest rank among all its neighbors (in  $C$ ), we call it a *root component*. Since every component except root components connects to a component with higher rank, the resulting structure is a set of disjoint rooted trees.

In the next step, we will merge all components of each tree into a single new component such that all vertices that are part of some component in this tree receive the label of the root. Consider a tree  $\mathcal{T}$ . We proceed level-wise (in parallel for all trees) and start the merging of components at the leaves that are connected to a (lower-ranking) parent component  $C$ .

LEMMA 2.5. *There is a distributed algorithm that merges all trees of the DRR forest in  $\tilde{O}(dn/k^2)$  rounds with high probability, where  $d$  is the largest depth of any tree.*

PROOF. We proceed in  $d$  iterations by merging the (current) leaf components with their parents in the tree. Thus it is sufficient to analyze the time complexity of a single iteration. To this end, we describe a procedure that changes the component labels of all vertices that are in leaf components in the DRR forest to the label of the respective parent in  $\tilde{O}(n/k^2)$  rounds.

At the beginning of each iteration, we select a new proxy for each component  $C$  by querying the shared hash function  $h_{j,\rho}(C)$ , where  $j$  is the index of the current phase and  $\rho$  is the index of the current iteration. This ensures that there are no dependencies between the proxies used in each iteration. We know from Lemma 2.4 that there is a message schedule such that leaf proxies can communicate with their respective parent proxy in  $\tilde{O}(n/k^2)$  rounds (w.h.p.) and vice versa, and thus every leaf proxy knows the component label of its parent. We have already shown in Lemma 2.3 that we can deliver a message from each component part to its respective proxy (when combining the sketches) in  $\tilde{O}(n/k^2)$  rounds. Hence, by re-running this message schedule, we can

<sup>9</sup>In place of DRR trees, an alternative and simpler idea is the following. Let every component select a number in  $[0, 1]$ . A merging can be done only if the outgoing edge connects a component with ID 0 to a component with ID 1. One can show that this merging procedure also gives the same time bound.

broadcast the parent label from the leaf proxy to each component part in the same time. Each machine that receives the parent label locally changes the component label of the vertices that are in the corresponding part.  $\square$

The following result is given in [8]. To keep the paper self-contained, in the Appendix we provide a more direct and simpler proof for this result.

LEMMA 2.6 ([8, THEOREM 11]). *The depth of each DRR tree is  $O(\log n)$  with high probability.*

## 2.6 Analysis of the Time Complexity

We now show that the number of phases required by the algorithm to determine the connected components of the input graph is  $O(\log n)$ . At the beginning of each phase  $i$ , distributed across the  $k$  machines there are  $c_i$  distinct components. At the beginning of the algorithm each node is identified as a component, and thus  $c_0 = n$ . The algorithm ends at the completion of phase  $\varphi$ , where  $\varphi$  is the smallest integer such that  $c_\varphi = cc(G)$ , where  $cc(G)$  denotes the number of connected components of the input graph  $G$ . If pairs of components were merged in each phase, it would be straightforward to show that the process would terminate in  $O(\log n)$  phases. However, in our algorithm each component connects to its neighboring component if and only if the latter has a higher rank. Nevertheless, it is not difficult to show that this slightly different process also terminates in  $O(\log n)$  phases w.h.p. (that is, components gets merged “often enough”). The intuition for this is that, since components’ ranks are taken randomly, for each component the probability that its neighboring component has a higher rank is exactly one half. Hence, on average half of the components will not be merged with their own neighbor: each of these components thus becomes a root of one component, which means that, on average, the number of new components is half as well.

LEMMA 2.7. *After  $12 \log n$  phases, the component labels of the vertices correspond to the connected components of  $G$  with high probability.*

PROOF. Replace the  $c_i$ ’s with corresponding random variables  $C_i$ ’s, and consider the stochastic process defined by the sequence  $C_0, C_1, \dots, C_\varphi$ . Let  $\bar{C}_i$  be the random variable that counts the number of components that actually participate at the merging process of phase  $i$ , because they do have an outgoing edge to another component. Call these components *participating components*. Clearly, by definition,  $\bar{C}_i \leq C_i$ .

We now show that, for every phase  $i \in [\varphi - 1]$ ,  $E[E[\bar{C}_{i+1} \mid \bar{C}_i]] \leq E[\bar{C}_i]/2$ . To this end, fix a generic phase  $i$  and a random ordering of its  $\bar{C}_i$  participating components. Define random variables  $X_{i,1}, X_{i,2}, \dots, X_{i,\bar{C}_i}$  where  $X_{i,j}$  takes value 1 if the  $j$ -th participating component will be a root of a participating tree/component for phase  $i + 1$ , and 0 otherwise. Then,  $\bar{C}_{i+1} \mid \bar{C}_i = \sum_{j=1}^{\bar{C}_i} X_{i,j}$  is the number of participating components for phase  $i + 1$ , assuming that there are  $C_i$  participating components for phase  $i$ . As we noticed before, for any  $i \in [\varphi - 1]$  and  $j \in [\bar{C}_i]$ , the probability that a participating component will not be merged to its neighboring component, and thus become a root of a tree/component for phase  $i + 1$  is exactly one half. Therefore,

$$\Pr(X_{i,j} = 1) \leq 1/2.$$

Hence, by the linearity of expectation, we have that

$$E[\bar{C}_{i+1} \mid \bar{C}_i] = \sum_{j=1}^{\bar{C}_i} E[X_{i,j}] = \sum_{j=1}^{\bar{C}_i} \Pr(X_{i,j} = 1) \leq \frac{\bar{C}_i}{2}.$$

Then, using again the linearity of expectation,

$$\mathbb{E}[\mathbb{E}[\bar{C}_{i+1} \mid \bar{C}_i]] \leq \mathbb{E}\left[\frac{\bar{C}_i}{2}\right] = \frac{\mathbb{E}[\bar{C}_i]}{2}.$$

We now leverage this result to prove the claimed statement. Let us call a phase *successful* if it reduces the number of participating components by a factor of at most  $3/4$ . By Markov's inequality, the probability that phase  $i$  is not successful is

$$\Pr\left(\mathbb{E}[\bar{C}_{i+1} \mid \bar{C}_i] > \frac{3}{4}\mathbb{E}[\bar{C}_i]\right) < \frac{\mathbb{E}[\mathbb{E}[\bar{C}_{i+1} \mid \bar{C}_i]]}{(3/4)\mathbb{E}[\bar{C}_i]} \leq \frac{\mathbb{E}[\bar{C}_i]}{2} \cdot \frac{4}{3\mathbb{E}[\bar{C}_i]} = \frac{2}{3},$$

and thus the probability that a phase of the algorithm is successful is at least  $1/3$ . Now consider a sequence of  $12 \log n$  phases of the algorithm. We shall prove that within that many phases the algorithm w.h.p. has reduced the number of participating components a sufficient number of times so that the algorithm has terminated, that is,  $\varphi \leq 12 \log n$  w.h.p. Let  $X_i$  be an indicator variable that takes value 1 if phase  $i$  is successful, and 0 otherwise (this also includes the case that the  $i$ -th phase does not take place because the algorithm already terminated). Let  $X = \sum_{i=1}^{12 \log n} X_i$  be the number of successful phases out of the at most  $12 \log n$  phases of the algorithm. Since  $\Pr(X_i = 1) \geq 1/3$ , by the linearity of expectation we have that

$$\mathbb{E}[X] = \sum_{i=1}^{12 \log n} \mathbb{E}[X_i] = \sum_{i=1}^{12 \log n} \Pr(X_i = 1) \geq \frac{12 \log n}{3} = 4 \log n.$$

As the  $X_i$ 's are independent we can apply a standard Chernoff bound, which gives

$$\Pr(X \leq \log n) \leq e^{-4 \log n (3/4)^2 / 2} = e^{-\frac{9}{8} \log n} < \frac{1}{n}.$$

Hence, with high probability  $12 \log n$  phases are enough to determine all the components of the input graph.  $\square$

**THEOREM 2.8.** *There exists a distributed algorithm for the  $k$ -machine model that determines the connected components of a graph  $G$  in  $\tilde{O}(n/k^2)$  rounds with high probability.*

**PROOF.** By Lemma 2.7, the algorithm finishes in  $O(\log n)$  phases with high probability. To analyze the time complexity of an individual phase, recall that it takes  $\tilde{O}(n/k^2)$  rounds to sample an outgoing edge (see Lemma 2.3). Then, building the DRR forest requires  $\tilde{O}(n/k^2)$  additional rounds, according to Lemma 2.4. Merging each DRR tree  $\mathcal{T}$  in a level-wise fashion (in parallel) takes  $\tilde{O}(dn/k^2)$  rounds (see Lemma 2.5), where  $d$  is the depth of  $\mathcal{T}$  which, by virtue of Lemma 2.6, is bounded by  $O(\log n)$ . Since each of these time bounds hold with high probability, and the algorithm consists of  $O(\log n)$  phases with high probability, by the union bound we conclude that the total time complexity of the algorithm is  $\tilde{O}(n/k^2)$  with high probability.  $\square$

We conclude the section by noticing that it is easy to output the actual number of connected components after the termination of the algorithm: every machine just needs to send, for each component label it holds, one message to the proxy machine of that component containing the component's label; then, such proxies locally count the number of distinct labels received, and send such a count to one predetermined machine. Since the communication is performed via the components' proxies, it follows from Lemma 2.1 that the first step takes  $\tilde{O}(n/k^2)$  rounds w.h.p., while the second step takes only one round.

### 3 APPLICATIONS

In this section we describe how to use our fast connectivity algorithm as a building block to solve several other fundamental graph problems in the  $k$ -machine model in time  $\tilde{O}(n/k^2)$ .

#### 3.1 Constructing a Minimum Spanning Tree

Given a weighted graph where each edge  $e = (u, v)$  has an associated weight  $w(e)$ , initially known to both the home machines of  $u$  and  $v$ , the minimum spanning tree (MST) problem asks to output a set of edges that form a tree, connect all nodes, and have the minimum possible total weight. Klauck et al. [24] show that  $\tilde{O}(n/k)$  rounds are necessary for constructing any spanning tree (ST), assuming that, for every spanning tree edge  $e = (u, v)$ , the home machine of  $u$  and the home machine of  $v$  must *both* output  $(u, v)$  as being part of the ST. Here we show that we can break the  $\tilde{O}(n/k)$  barrier, under the slightly less stringent requirement that each spanning tree edge  $e$  is returned by at least *one* machine, but not necessarily by both the home machines of  $u$  and  $v$ .

Our algorithm mimics the multi-pass MST construction procedure of [1], originally devised for the (centralized) streaming model. To this end we modify our connectivity procedure of Section 2, by ensuring that when a component proxy  $C$  chooses an outgoing edge  $e$ , this is the *minimum-weight outgoing edge* (MWOE) of  $C$  with high probability.

We now describe the  $i$ -th phase of this MST construction in more detail. Analogously to the connectivity algorithm, the proxy of each component  $C$  determines an outgoing edge  $e_0$  which, by the guarantees of our sketch construction (Lemma 2.2), is chosen uniformly at random from all possible outgoing edges of  $C$ .

We then repeat the following edge-elimination process  $t = \Theta(\log n)$  times: The proxy broadcasts  $w(e_0)$  to every component part of  $C$ . By Lemma 2.3 this communication is possible in  $\tilde{O}(n/k^2)$  rounds. Upon receiving this message, the machine  $M$  of a part  $P$  of  $C$  constructs a new sketch  $\mathbf{s}_u$  for each  $u \in P$ , but first zeroes out all entries in  $\mathbf{a}_u$  that refer to edges of weight  $> w(e_0)$ . (See Section 2.3 for a more detailed description of  $\mathbf{a}_u$  and  $\mathbf{s}_u$ .) Again, we combine the sketches of all vertices of all parts of  $C$  at the proxy of  $C$ , which in turn samples a new outgoing edge  $e_1$  for  $C$ . Since each time we sample a randomly chosen edge and eliminate all higher weight edges, it is easy to see that the edge  $e_t$  is the MWOE of  $C$  w.h.p. Thus, the proxy machine of  $C$  includes the edge  $e_t$  as part of the MST output. Note that this additional elimination procedure incurs only a logarithmic time complexity overhead.

At the end of each phase, we proceed by (virtually) merging the components along their MWOEs in a similar manner as for the connectivity algorithm (see Section 2.5), thus requiring  $\tilde{O}(n/k^2)$  rounds in total.

Let  $E$  be the set of added outgoing edges. Since the components of the connectivity algorithm eventually match the actual components of the input graph, the graph  $H$  on the vertices  $V(G)$  induced by  $E$  connects all vertices of  $G$ . Moreover, since components are merged according to the trees of the DRR-process (see Section 2.5), it follows that  $H$  is cycle-free.

We can now fully classify the complexity of the MST problem in the  $k$ -machine model.

**THEOREM 3.1.** *There exists a distributed algorithm for the  $k$ -machine model that outputs an MST in*

- $\tilde{O}(n/k^2)$  rounds, if each MST-edge is output by at least one machine, or in
- $\tilde{O}(n/k)$  rounds, if each MST-edge  $e$  is output by both machines that hold an endpoint of  $e$ .

*Both bounds are tight up to polylogarithmic factors.*

#### 3.2 Approximate Min-Cut

Here we show the following result for the min-cut problem in the  $k$ -machine model.

**THEOREM 3.2.** *There exists an  $O(\log n)$ -approximation algorithm for the min-cut problem in the  $k$ -machine model that runs in  $\tilde{O}(n/k^2)$  rounds with high probability.*

**PROOF.** We use exponentially growing sampling probabilities for sampling edges and then check connectivity, leveraging a result by [20]. This procedure was derived in [16] for the CONGEST model, and can be implemented in the  $k$ -machine model as well, where we use our fast connectivity algorithm (in place of Thurimella's algorithm [47] used in [16]). The time complexity is dominated by the connectivity-testing procedure, and thus is  $\tilde{O}(n/k^2)$  w.h.p.  $\square$

### 3.3 Algorithms for Graph Verification Problems

It is well known that graph connectivity is an important building block for several graph verification problems (see, e.g., [12]). We now analyze some of such problems, formally defined, e.g., in Section 2.4 of [12], in the  $k$ -machine model.

**THEOREM 3.3.** *There exist algorithms for the  $k$ -machine model that solve the following verification problems in  $\tilde{O}(n/k^2)$  rounds with high probability: spanning connected subgraph, cycle containment,  $e$ -cycle containment, cut,  $s$ - $t$  connectivity, edge on all paths,  $s$ - $t$  cut, bipartiteness.*

**PROOF.** We discuss each problem separately.

**Cut verification** Remove the edges of the given cut from  $G$ , and then check whether the resulting graph is connected.

**$s$ - $t$  connectivity verification** Run the connectivity algorithm and then verify whether  $s$  and  $t$  are in the same connected component by checking whether they have the same label.

**Edge on all paths verification** Since  $e$  lies on all paths between  $u$  and  $v$  iff  $u$  and  $v$  are disconnected in  $G \setminus \{e\}$ , we can simply use the  $s$ - $t$  connectivity verification algorithm of previous point.

**$s$ - $t$  cut verification** To verify if a subgraph is an  $s$ - $t$  cut, simply verify  $s$ - $t$  connectivity of the graph after removing the edges of the subgraph.

**Bipartiteness verification** Use the connectivity algorithm and the reduction presented in Section 3.3 of [1].

**Spanning connected subgraph, cycle containment, and  $e$ -cycle containment verification** These also follow from the reductions given in [12].  $\square$

## 4 LOWER BOUNDS FOR VERIFICATION PROBLEMS

In this section we show that  $\tilde{\Omega}(n/k^2)$  rounds is a fundamental lower bound for many graph verification problems in the  $k$ -machine model. To this end we will use results from the classical theory of communication complexity [25], a popular way to derive lower bounds in distributed message-passing models [12, 38, 39].

Even though many verification problems are known to satisfy a lower bound of  $\tilde{\Omega}(D + \sqrt{n})$  in the classic distributed CONGEST model [12], the reduction of [12] encodes a  $\Theta(\sqrt{n})$ -instance of set disjointness, requiring at least one node to receive  $\tilde{\Theta}(\sqrt{n})$  information across a single short "highway" path or via  $\Theta(\sqrt{n})$  longer paths of length  $\Theta(\sqrt{n})$ . Moreover, we assume the random vertex partition model, whereas the results of [12] assume a worst case distribution. Lastly, any pair of machines can communicate directly in the  $k$ -machine model, thus breaking the  $\Omega(D)$  bound for the CONGEST model.

Our complexity bounds follow from the communication complexity of 2-player *set disjointness* in the *random input partition model* (see [24]). While in the standard model of communication complexity there are 2 players, Alice and Bob, and Alice (resp., Bob) receives an input vector  $X$  (resp.,  $Y$ ) of  $b$  bits [25], in the random input partition model Alice receives  $X$  and, in addition, each

bit of  $Y$  has probability  $1/2$  to be revealed to Alice. Bob's input is defined similarly with respect to  $X$ . In the set disjointness problem, Alice and Bob must output 1 if and only if there is no index  $i$  such that  $X[i] = Y[i] = 1$ . The following result holds.

LEMMA 4.1 ([24, LEMMA 3.2]). *For some constant  $\epsilon > 0$ , every randomized communication protocol that solves set disjointness in the random input partition model of 2-party communication complexity with probability at least  $1 - \epsilon$ , requires  $\Omega(b)$  bits.*

Now we can show the main result of this section.

THEOREM 4.2. *There exists a constant  $\gamma > 0$  such that any  $\gamma$ -error algorithm  $\mathcal{A}$  has round complexity of  $\tilde{\Omega}(n/k^2)$  on an  $n$ -vertex graph of diameter 2 in the  $k$ -machine model, if  $\mathcal{A}$  solves any of the following problems: connectivity, spanning connected subgraph, cycle containment,  $e$ -cycle containment,  $s$ - $t$ -connectivity, cut, edge on all paths, and  $s$ - $t$ -cut.*

PROOF. The high-level idea of the proof is similar to the simulation theorem of [12]. We present the argument for the spanning connected subgraph (SCS) problem, defined next. Following [12], similar arguments can be established for each of the other problems.

In the spanning connected subgraph problem we are given a graph  $G$  and a subgraph  $H \subseteq G$  and we want to verify whether  $H$  spans  $G$  and is connected. We will show, through a reduction from 2-party set disjointness, that any algorithm for SCS in the  $k$ -machine model requires  $\tilde{\Omega}(n/k^2)$  rounds.

Given an instance of the 2-party set disjointness problem in the random partition model we will construct the following input graphs  $G$  and  $H$ . The nodes of  $G$  consist of 2 special nodes  $s$  and  $t$ , and nodes  $u_1, \dots, u_b, v_1, \dots, v_b$ , for  $b = (n - 2)/2$ . (For clarity of presentation, we assume that  $(n - 2)/2$  and  $k/2$  are integers.) The edges of  $G$  consist of the edges  $(s, t)$ ,  $(u_i, v_i)$ ,  $(s, u_i)$ ,  $(v_i, t)$ , for  $1 \leq i \leq b$ .

Let  $\mathcal{M}_A$  be the set of machines simulated by Alice, and let  $\mathcal{M}_B$  be the set of machines simulated by Bob, where  $|\mathcal{M}_A| = |\mathcal{M}_B| = k/2$ . First, Alice and Bob use shared randomness to choose the machines  $M_X$  and  $M_Y$  that receive the vertices  $s$  and  $t$ . If  $M_X \neq M_Y$ , then Alice assigns  $t$  to a machine chosen randomly from  $\mathcal{M}_A$ , and Bob assigns  $s$  to a random machine in  $\mathcal{M}_B$ . Otherwise, if  $M_X$  and  $M_Y$  denote the same machine, Alice and Bob output 0 and terminate the simulation.

The subgraph  $H$  is determined by the disjointness input vectors  $X$  and  $Y$  as follows:  $H$  contains all nodes of  $G$  and the edges  $(u_i, v_i)$ ,  $(s, t)$ ,  $1 \leq i \leq b$ . Recall that, in the random partition model,  $X$  and  $Y$  are randomly distributed between Alice and Bob, but Alice knows all  $X$  and Bob knows all of  $Y$ . Hence, Alice and Bob mark the corresponding edges as being part of  $H$  according to their respective input bits. That is, if Alice received  $X[i]$  (i.e. Bob did *not* receive  $X[i]$ ), she assigns the node  $u_i$  to a random machine in  $\mathcal{M}_A$  and adds the edge  $(s, u_i)$  to  $H$  if and only if  $X[i] = 0$ . Similarly, the edge  $(v_i, t)$  is added to  $H$  if and only if  $Y[i] = 0$  (by either Alice or Bob depending on who receives  $Y[i]$ ). See Figure 1. Note that, since  $X$  and  $Y$  were assigned according to the random input partition model, the resulting distribution of vertices to machines adheres to the random vertex partition model. Clearly,  $H$  is an SCS if and only if  $X$  and  $Y$  are disjoint.

We describe the simulation from Alice's point of view (the simulation for Bob is similar): Alice locally maintains a counter  $r_A$ , initialized to 1, that represents the current round number. Then, she simulates the run of  $\mathcal{A}$  on each of her  $k/2$  machines, yielding a set of  $\ell$  messages  $m_1, \dots, m_\ell$  of  $O(\text{polylog}(n))$  bits each that need to be sent to Bob to simulate the algorithm on his machines in the next round. By construction, we have that  $0 \leq \ell \leq \lceil k^2/4 \rceil$ . To send these messages in the (asynchronous) 2-party random partition model of communication complexity, Alice sends a message  $\langle \ell, (M_1, m_1, M_2), \dots, (M_\ell, m_\ell, M_{\ell+1}) \rangle$  to Bob, where a tuple  $(M_i, m_i, M_{i+1})$  corresponds to a message  $m_i$  generated by machine  $M_i$  simulated by Alice and destined to machine  $M_{i+1}$  simulated

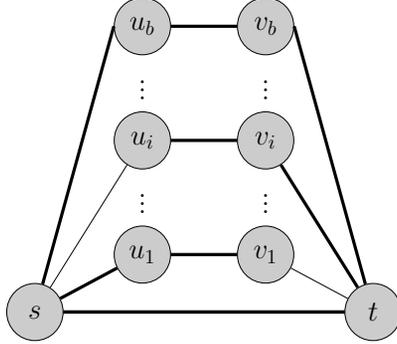


Fig. 1. The graph construction for the spanning connected subgraph problem, given a set disjointness instance where  $X[1] = 0$ ,  $Y[1] = 1$ ,  $X[i] = 1$ ,  $Y[i] = 0$ , and  $X[b] = Y[b] = 0$ . The thick edges are the edges of subgraph  $H$ . The subgraph  $H$  contains all edges  $(u_i, v_i)$  ( $1 \leq i \leq b$ ) and  $(s, t)$ ; the remaining edges of  $H$  are determined by the input vectors  $X$  and  $Y$  of the set disjointness instance.

at Bob. Upon receiving this message, Bob increases its own round counter and then locally simulates the next round of his machines by delivering the messages to the appropriate machines. Adding the source and destination fields to each message incurs an overhead of only  $O(\log k) = O(\log n)$  bits, hence the total communication generated by simulating a single round of  $\mathcal{A}$  is upper bounded by  $\tilde{O}(k^2)$ . Therefore, if  $\mathcal{A}$  takes  $T$  rounds to solve SCS in the  $k$ -machine model, then this gives us an  $O(Tk^2 \text{polylog}(n))$ -bit communication complexity protocol for set disjointness in the random partition model, as the communication between Alice and Bob is determined by the communication across the  $\Theta(k^2)$  links required for the simulation, each of which can carry  $O(\text{polylog}(n))$  bits per round. Note that if  $\mathcal{A}$  errs with probability at most  $\gamma$ , then the simulation errs with probability at most  $\gamma + 1/k$ , where the extra  $1/k$  term comes from the possibility that machines  $M_X$  and  $M_Y$  refer to the same machine. For large enough  $k$  and small enough  $\gamma$  we have  $\gamma + 1/k < \epsilon$ . It follows that we need to simulate at least  $T = \tilde{\Omega}(n/k^2)$  many rounds, since by Lemma 4.1 the set disjointness problem requires  $\Omega(b)$  bits in the random partition model, when the error is smaller than  $\epsilon$ .  $\square$

Interestingly, our lower bounds hold even for graphs of diameter 2, which is in contrast to the analogous results for the classic distributed CONGEST model assumed in [12]. We remark that the lower bound for connectivity verification was already shown in [24].

## 5 CONCLUSIONS

There are several natural directions for future work. Our connectivity algorithm is randomized: it would be interesting to study the deterministic complexity of graph connectivity in the  $k$ -machine model. Specifically, does graph connectivity admit a  $\tilde{O}(n/k^2)$  deterministic algorithm? Investigating higher-order connectivity, such as 2-edge/vertex connectivity, is also an interesting research direction. A general question motivated by the algorithms presented in this paper is whether one can design algorithms that have superlinear scaling in  $k$  for other fundamental graph problems. Some recent results in this directions are in [42]. When designing algorithms for real-world graph processing systems, besides minimizing the round complexity, it is also important that the total number of messages exchanged be kept as low as possible; this direction, not new in the field of distributed computing (see, e.g., [41]), is also worth pursuing.

## APPENDIX

## A.1 Proof of Lemma 2.6

PROOF. Consider one phase of the algorithm, and suppose that during that phase there are  $n$  components. (In one phase there are  $c \leq n$  components, thus setting  $c = n$  gives a valid upper bound to the height of each DRR tree in that phase.) Each component picks a random rank from  $[0, 1]$ . Thus, all ranks are distinct with high probability. If the target component's rank is higher, then the source component connects to it, otherwise the source component becomes a root of a DRR tree.

Consider an arbitrary component of the graph, and consider the (unique) path  $P$  starting from the node that represents the component to the root of the tree that contains it. Let  $|P|$  be the number of nodes of  $P$ , and assign indexes to the  $|P|$  nodes of  $P$  according to their position in the path from the selected node to the root. (See Figure 2.)

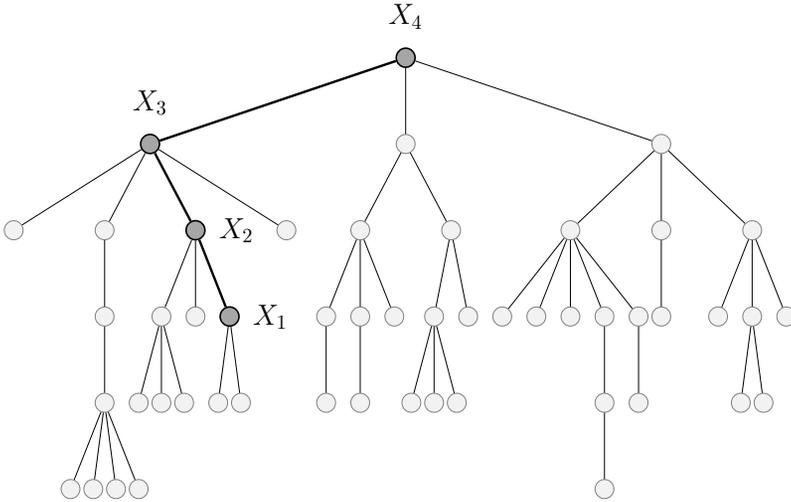


Fig. 2. One DRR tree, and one path from one node to the root of the tree. Nodes of the path are labeled with the indicator variable associated to them, indexed by the position of the node in the path.

For each  $i \in [|P|]$ , define  $X_i$  as the indicator variable that takes value 1 if node  $i$  is not the root of  $P$ , and 0 otherwise. Then,  $X = \sum_{i=1}^{|P|} X_i$  is the length of the path  $P$ . Because of the random choice for the outgoing edge made by components' parts, the outgoing edge of each component is to a random (and distinct) component. This means that, for each  $j \leq |P|$ , the ranks of the first  $j$  nodes of the path form a set of  $j$  random values in  $[0, 1]$ . Hence, the probability that a new random value in  $[0, 1]$  is higher than the rank of the  $j$ -th node of the path is the probability that the new random value is higher than all the  $j$  previously chosen random values (that is, the probability its value is the highest among all the first  $j$  values of the path), and this probability is at most  $1/(j + 1)$ . Thus,  $\Pr(X_i = 1) \leq 1/(i + 1)$ . Hence, by the linearity of expectation, the expected height of a path in a

tree produced by the DRR procedure is

$$E[X] = \sum_{i=1}^{|P|} E[X_i] \leq \sum_{i=1}^n E[X_i] = \sum_{i=1}^n \Pr(X_i = 1) \leq \sum_{i=1}^n \frac{1}{i+1} \leq \log(n+1).$$

Notice that the  $X_i$ 's are independent (but not identically distributed) random variables, since the probability that the  $i$ -th smallest ranked node is not a root depends only on the random neighbor that it picks, and is independent of the choices of the other nodes. Thus, applying a standard Chernoff bound (see, e.g., [35]) we have

$$\Pr(X \geq 6 \log(n+1)) \leq 2^{-6 \log(n+1)} = \frac{1}{(n+1)^6}.$$

Applying the union bound over all the at most  $n$  paths concludes the proof.  $\square$

## ACKNOWLEDGMENTS

The authors would like to thank Mohsen Ghaffari, Seth Gilbert, Andrew McGregor, Danupon Nanongkai, and Sriram V. Pemmaraju for helpful discussions.

## REFERENCES

- [1] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. 2012. Analyzing graph structure via linear measurements. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 459–467.
- [2] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. 2012. Graph sketches: sparsification, spanners, and subgraphs. In *Proceedings of the 31st ACM Symposium on Principles of Database Systems (PODS)*. 5–14.
- [3] Noga Alon, László Babai, and Alon Itai. 1986. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms* 7, 4 (1986), 567–583.
- [4] Noga Alon, Ronitt Rubinfeld, Shai Vardi, and Ning Xie. 2012. Space-efficient local computation algorithms. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1132–1139.
- [5] Sayan Bandyapadhyay, Tanmay Inamdar, Shreyas Pai, and Sriram V. Pemmaraju. 2018. Near-optimal clustering in the  $k$ -machine model. In *Proceedings of the 19th International Conference on Distributed Computing and Networking (ICDCN)*.
- [6] Otakar Borůvka. 1926. O jistém problému minimálním (about a certain minimal problem). *Práce Mor. Přírodoved. Spol. v Brně III* 3 (1926).
- [7] Keren Censor-Hillel, Petteri Kaski, Janne H. Korhonen, Christoph Lenzen, Ami Paz, and Jukka Suomela. 2015. Algebraic methods in the congested clique. In *Proceedings of the 34th ACM Symposium on Principles of Distributed Computing (PODC)*. 143–152.
- [8] Jen-Yeu Chen and Gopal Pandurangan. 2012. Almost-optimal gossip-based aggregate computation. *SIAM J. Comput.* 41, 3 (2012), 455–483.
- [9] Avery Ching, Sergey Edunov, Maja Kabiljo, Dionysios Logothetis, and Sambavi Muthukrishnan. 2015. One trillion edges: graph processing at Facebook-scale. *Proc. VLDB Endow.* 8, 12 (2015), 1804–1815.
- [10] Fan Chung and Olivia Simpson. 2015. Distributed algorithms for finding local clusters using heat kernel Pagerank. In *Proceedings of the 12th Workshop on Algorithms and Models for the Web-graph (WAW)*. 77–189.
- [11] Graham Cormode and Donatella Firmani. 2014. A unifying framework for  $\ell_0$ -sampling algorithms. *Distributed and Parallel Databases* 32, 3 (2014), 315–335.
- [12] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. 2012. Distributed verification and hardness of distributed approximation. *SIAM J. Comput.* 41, 5 (2012), 1235–1265.
- [13] Andrew Drucker, Fabian Kuhn, and Rotem Oshman. 2014. On the power of the congested clique model. In *Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing (PODC)*. 367–376.
- [14] Michael Elkin, Hartmut Klauck, Danupon Nanongkai, and Gopal Pandurangan. 2014. Can quantum communication speed up distributed computation?. In *Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing (PODC)*. 166–175.
- [15] Robert G. Gallager, Pierre A. Humblet, and Philip M. Spira. 1983. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Program. Lang. Syst.* 5, 1 (1983), 66–77.
- [16] Mohsen Ghaffari and Fabian Kuhn. 2013. Distributed minimum cut approximation. In *Proceedings of the 27th International Symposium on Distributed Computing (DISC)*. 1–15.

- [17] Mohsen Ghaffari and Merav Parter. 2016. MST in log-star rounds of congested clique. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC)*. 19–28.
- [18] James W. Hegeman, Gopal Pandurangan, Sriram V. Pemmaraju, Vivek B. Sardeshmukh, and Michele Scquizzato. 2015. Toward optimal bounds in the congested clique: graph connectivity and MST. In *Proceedings of the 34th ACM Symposium on Principles of Distributed Computing (PODC)*. 91–100.
- [19] Hossein Jowhari, Mert Saglam, and Gábor Tardos. 2011. Tight bounds for  $L_p$  samplers, finding duplicates in streams, and related problems. In *Proceedings of the 30th ACM Symposium on Principles of Database Systems (PODS)*. 49–58.
- [20] David R. Karger. 1994. Random sampling in cut, flow, and network design problems. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing (STOC)*. 648–657.
- [21] David R. Karger, Philip N. Klein, and Robert E. Tarjan. 1995. A randomized linear-time algorithm to find minimum spanning trees. *J. ACM* 42, 2 (1995), 321–328.
- [22] Howard J. Karloff, Siddharth Suri, and Sergei Vassilvitskii. 2010. A model of computation for MapReduce. In *Proceedings of the 21st annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 938–948.
- [23] Valerie King, Shay Kutten, and Mikkel Thorup. 2015. Construction and impromptu repair of an MST in a distributed network with  $o(m)$  communication. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing (PODC)*. 71–80.
- [24] Hartmut Klauck, Danupon Nanongkai, Gopal Pandurangan, and Peter Robinson. 2015. Distributed computation of large-scale graph problems. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 391–410.
- [25] Eyal Kushilevitz and Noam Nisan. 1997. *Communication Complexity*. Cambridge University Press.
- [26] Shay Kutten, Gopal Pandurangan, David Peleg, Peter Robinson, and Amitabh Trehan. 2015. Sublinear bounds for randomized leader election. *Theoret. Comput. Sci.* 561 (2015), 134–143.
- [27] Silvio Lattanzi, Benjamin Moseley, Siddharth Suri, and Sergei Vassilvitskii. 2011. Filtering: a method for solving graph problems in MapReduce. In *Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 85–94.
- [28] Christoph Lenzen. 2013. Optimal deterministic routing and sorting on the congested clique. In *Proceedings of the 32nd ACM Symposium on Principles of Distributed Computing (PODC)*. 42–50.
- [29] Christoph Lenzen and Roger Wattenhofer. 2016. Tight bounds for parallel randomized load balancing. *Distrib. Comput.* 29, 2 (2016), 127–142.
- [30] Jure Leskovec, Anand Rajaraman, and Jeffrey David Ullman. 2014. *Mining of Massive Datasets*. Cambridge University Press.
- [31] Zvi Lotker, Boaz Patt-Shamir, Elan Pavlov, and David Peleg. 2005. Minimum-weight spanning tree construction in  $O(\log \log n)$  communication rounds. *SIAM J. Comput.* 35, 1 (2005), 120–131.
- [32] Nancy A. Lynch. 1996. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc.
- [33] Grzegorz Malewicz, Matthew H. Austern, Aart J. C. Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. 2010. Pregel: a system for large-scale graph processing. In *Proceedings of the 2010 ACM International Conference on Management of Data (SIGMOD)*. 135–146.
- [34] Andrew McGregor. 2014. Graph stream algorithms: a survey. *SIGMOD Record* 43, 1 (2014), 9–20.
- [35] Michael Mitzenmacher and Eli Upfal. 2005. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press.
- [36] Danupon Nanongkai. 2014. Distributed approximation algorithms for weighted shortest paths. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*. 565–573.
- [37] Danupon Nanongkai, Atish Das Sarma, and Gopal Pandurangan. 2011. A tight unconditional lower bound on distributed randomwalk computation. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing (PODC)*. 257–266.
- [38] Rotem Oshman. 2014. Communication complexity lower bounds in distributed message-passing. In *Proceedings of the 21th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*. 14–17.
- [39] Gopal Pandurangan, David Peleg, and Michele Scquizzato. 2016. Message lower bounds via efficient network synchronization. In *Proceedings of the 23rd International Colloquium on Structural Information and Communication Complexity (SIROCCO)*. 75–91.
- [40] Gopal Pandurangan, Peter Robinson, and Michele Scquizzato. 2016. Fast distributed algorithms for connectivity and MST in large graphs. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 429–438.
- [41] Gopal Pandurangan, Peter Robinson, and Michele Scquizzato. 2017. A time- and message-optimal distributed algorithm for minimum spanning trees. In *Proceedings of the 49th Annual ACM Symposium on the Theory of Computing (STOC)*. 743–756.

- [42] Gopal Pandurangan, Peter Robinson, and Michele Scquizzato. 2018. On the distributed complexity of large-scale graph computations. In *Proceedings of the 30th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. To appear.
- [43] David Peleg. 2000. *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial and Applied Mathematics.
- [44] Sriram V. Pemmaraju and Vivek B. Sardeshmukh. 2016. Super-fast MST algorithms in the congested clique using  $o(m)$  messages. In *Proceedings of the 36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*. 47:1–47:15.
- [45] Judy Qiu, Shantenu Jha, Andre Luckow, and Geoffrey C. Fox. 2014. Towards HPC-ABDS: an initial high-performance big data stack. Available: <http://grids.ucs.indiana.edu/ptliupages/publications/nist-hpc-abds.pdf>.
- [46] Isabelle Stanton. 2014. Streaming balanced graph partitioning algorithms for random graphs. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1287–1301.
- [47] Ramakrishna Thurimella. 1997. Sub-linear distributed algorithms for sparse certificates and biconnected components. *J. Algorithms* 23, 1 (1997), 160–179.
- [48] Yuanyuan Tian, Andrey Balmin, Severin Andreas Corsten, Shirish Tatikonda, and John McPherson. 2013. From “Think Like a Vertex” to “Think Like a Graph”. *PVLDB* 7, 3 (2013), 193–204.
- [49] Leslie G. Valiant. 1982. A scheme for fast parallel communication. *SIAM J. Comput.* 11, 2 (1982), 350–361.
- [50] Leslie G. Valiant. 1990. A bridging model for parallel computation. *Commun. ACM* 33, 8 (1990), 103–111.
- [51] Sergei Vassilvitskii. 2015. Models for Parallel Computation (A Hitchhikers’ Guide to Massively Parallel Universes). (2015). <http://grigory.us/blog/massively-parallel-universes/>.
- [52] David P. Woodruff and Qin Zhang. 2017. When distributed computation is communication expensive. *Distrib. Comput.* 30, 5 (2017), 309–323.