Massively Parallel Computation via Remote Memory Access

SOHEIL BEHNEZHAD, University of Maryland

LAXMAN DHULIPALA, Carnegie Mellon University

HOSSEIN ESFANDIARI, JAKUB ŁĄCKI, VAHAB MIRROKNI, and WARREN SCHUDY,
Google Research

We introduce the *Adaptive Massively Parallel Computation* (AMPC) model, which is an extension of the *Massively Parallel Computation* (MPC) model. At a high level, the AMPC model strengthens the MPC model by storing all messages sent within a round in a distributed data store. In the following round, all machines are provided with random read access to the data store, subject to the same constraints on the total amount of communication as in the MPC model. Our model is inspired by the previous empirical studies of distributed graph algorithms [8, 30] using MapReduce and a distributed hash table service [17].

This extension allows us to give new graph algorithms with much lower round complexities compared to the best-known solutions in the MPC model. In particular, in the AMPC model we show how to solve maximal independent set in $O(1)$ rounds and connectivity/minimum spanning tree in $O(\log \log_{m/n} n)$ rounds both using $O(n^\delta)$ space per machine for constant $\delta < 1$. In the same memory regime for MPC, the best-known algorithms for these problems require $\text{poly log } n$ rounds. Our results imply that the 2-CYCLE conjecture, which is widely believed to hold in the MPC model, does not hold in the AMPC model.

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

Additional Key Words and Phrases: Datasets, neural networks, gaze detection, text tagging

ACM Reference format:

Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Łącki, Vahab Mirrokni, and Warren Schudy. 2021. Massively Parallel Computation via Remote Memory Access. *ACM Trans. Parallel Comput.* 8, 3, Article 13 (September 2021), 25 pages.

<https://doi.org/10.1145/3470631>

1 INTRODUCTION

The MapReduce computation framework was introduced in 2004 by Dean and Ghemawat [21]. It provides an easy to use platform for distributed computing, which handles synchronization and fault tolerance entirely on the framework side. Since its introduction MapReduce has inspired a number of other distributed computation frameworks, for example Spark [48], Hadoop [5], Flume-Java [16], Beam [3], Pregel [32], and Giraph [4].

Authors' addresses: S. Behnezhad, University of Maryland, 8125 Paint Branch Dr, College Park, MD 20742; L. Dhulipala, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213; H. Esfandiari, J. Łącki, V. Mirrokni, and W. Schudy, Google Research, 111 8th Ave, New York, NY 10011.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Association for Computing Machinery.

2329-4949/2021/09-ART13 \$15.00

<https://doi.org/10.1145/3470631>

Problem	AMPC	MPC
Connectivity	$O(\log \log_{m/n} n)$	$O(\log D + \log \log_{m/n} n)$ [10]
Minimum spanning tree	$O(\log \log_{m/n} n)$	$O(\log n)$
2-edge connectivity	$O(\log \log_{m/n} n)$	$O(\log D \cdot \log \log_{m/n} n)$ [2]
Maximal independent set	$O(1)$	$\tilde{O}(\sqrt{\log n})$ [25]
2-CYCLE	$O(1)$	$O(\log n)$
Forest Connectivity	$O(1)$	$O(\log D \cdot \log \log_{m/n} n)$ [2]

Fig. 1. Round complexities of our algorithms in the AMPC model compared to the state-of-the-art MPC algorithms. D denotes the diameter of the input graph. We consider the setting where the space per machine is sublinear in the number of vertices, that is $S = O(n^\epsilon)$ for constant $\epsilon < 1$.

These frameworks share many similarities, especially from a theoretical point of view and the **Massively Parallel Computation (MPC)** model [9, 26, 29] is widely accepted as the standard theoretical model of such systems. Due to the generality of the MPC model, it does not account for specific features used by emerging data-processing systems, such as **remote direct-memory access (RDMA)**s and **Remote procedure call (RPC)**s whose use may result in simpler and faster algorithms. In particular, in the context of graph algorithms, the major shortcoming of the MPC model is the fact that exploring the neighborhood of a vertex is costly. The widely believed 2-CYCLE conjecture [10, 24, 46] states that distinguishing between a graph being a cycle of length n from a graph consisting of two cycles of length $n/2$ requires $\Omega(\log n)$ rounds. Intuitively, finding even a single vertex at a distance k from a given vertex seems to require $\Omega(\log k)$ rounds. This restriction turns out to be, perhaps, *the* most important bottleneck in designing efficient MPC algorithms for graph problems when the space per machine is much smaller than n , the number of vertices. In practice, however, the local neighborhood of a vertex can be explored much more efficiently. In a broader context, a data stored on a remote machine can be read with only a few microsecond latency using hardware support for RDMA without requiring a synchronized round of communication.

In this article, we attempt to close this gap between theory and practice by introducing an extension of the MPC model and exploring its capabilities.

1.1 Our contribution

In this article, we introduce the **Adaptive Massively Parallel Computation (AMPC)** model, which is an extension of the MPC model. Roughly speaking, the AMPC model extends the MPC model by allowing machines to access a shared read-only memory *within a round*. We model this by assuming that all messages sent in a single round are written to a distributed data storage, which all machines can read from within the next round.

We demonstrate the usefulness of the new model by giving new algorithms for a number of fundamental graph problems. The corresponding round complexities are significantly lower than the complexities of the best-known algorithms in the MPC model. Our results are summarized in Figure 1. We highlight that our results imply that the 2-CYCLE conjecture does not hold in the AMPC model, separating it from the MPC model (under the conjecture).

Our model is inspired by two previous papers that combine the use of **distributed hash table service (DHT)** and MapReduce [8, 30]. The first paper [30] considers the problem of finding connected components. It shows that by using a DHT in addition to MapReduce, the bound on the number of rounds can be decreased from $O(\log^2 n)$ to $O(\log n \log \log n)$. It also demonstrates empirically that using a DHT can significantly decrease the number of rounds and the running time of the algorithm. The second paper [8] studies the distributed hierarchical clustering problem and shows how the use of a DHT facilitates the affinity clustering algorithm to scale up to graphs with trillions of edges.

However, neither of these papers defined a formal model of computation in which their algorithms operate. Our article is thus the first attempt to model the use of a distributed hash table along with a MapReduce-like distributed processing framework.

1.2 Organization of the Paper

In Section 2, we describe our model and discuss its properties. In particular, we discuss various practical aspects of implementing our new model. Then, in Section 3, we introduce notation that is used in the following sections. The remaining sections present our algorithms. We start with a very simple algorithm for solving the 2-CYCLE problem in Section 4. This is followed by algorithms for maximal independent set (Section 5), connectivity (Section 6), minimum spanning forest (Section 7), and forest connectivity (Section 8).

2 THE ADAPTIVE MASSIVELY PARALLEL COMPUTATION (AMPC) MODEL

In the AMPC model, we have an input of size N^1 , which is processed by P machines, each with space S . We make the standard assumption that the word size of the machines is $O(\log N)$. We denote by T the total space of all machines, i.e., $T = S \cdot P$. We assume that $S = \Theta(N^{1-\Omega(1)})$ and require that the total space is close to the input size, i.e., $T = O(N \text{ polylog } N)$.

In addition, in the AMPC model there is collection of distributed key-value stores, or DHT that we denote by $\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2, \dots$. Each DHT stores a collection of key-value pairs. When the DHT is queried for a given key, it returns all corresponding values associated with the key. Querying for a key that does not occur in the DHT results in an empty response. We require that both keys and values have constant size (i.e., both consist of a constant number of machine words). The input data is stored in $\mathcal{D}_0$ and uses a set of keys known to all machines (e.g., consecutive integers).

The computation consists of *rounds*. In the i th round, each machine can read data from $\mathcal{D}_{i-1}$ and write to $\mathcal{D}_i$. Within a round, each machine can make a total number of reads (henceforth called *queries*) and receive responses such that the sum of the number of queries and values returned by the responses is S . In addition, within a round a machine can perform S *writes* to $\mathcal{D}_i$, and perform arbitrary computation.² The total *communication* of a machine is equal to the total number of queries, values returned, and writes.

Querying for individual values. For the sake of simplicity, we will use the following computationally equivalent, and conceptually simpler model in the remainder of this article. In the simplified model, if $k > 1$ key-value pairs stored in the DHT have the same key x , we allow individual values to be accessed by querying for keys $(x, 1), \dots, (x, k)$. Note that the indices from 1 to k are assigned arbitrarily. Furthermore, the total number of writes to, or equivalently values associated with a given key can be at most S .

These models are equivalent (up to constant factors in the round complexity) for the following reason: In the original model, the fact that all values for a given key must be returned when querying that key immediately upper bounds the number of values associated with the key to S . To simulate the original model with the simpler one, after each round of computation, we can perform an additional round that maps over all keys and all values for a key, and emits pairs where the keys now store both the key and index, and the value is the value corresponding to this index. This transformation allows algorithms to query for individual values while only increasing the round complexity by a constant factor. The other direction simply performs multiple queries for a given key to fetch all (at most S) values associated with it.

¹The input size N should not be confused with the number of vertices in a graph, which we denote by n .

²The algorithms described in this article perform near-linear computation in each step (w.r.t. their available space). However, we do not put a bound on the amount of computation on a machine, as this is a usual assumption of the MPC model.

Adaptivity. The key property of the model is that the queries that a machine makes in a round may depend on the results of the previous queries it made in the same round. This is why we call our model *adaptive*. In particular, if g is a function from X to X and for each $x \in X$, then $\mathcal{D}_{i-1}$ stores a key-value pair $(x, g(x))$, then in round i a machine can compute $g^{(k)}(y)$ (here $g^{(1)}(x) = g(x)$ and $g^i(x) = g(g^{(i-1)}(x))$ for any $i > 1$) in a single round, provided that $k = O(S)$.

Comparison/Simulations. It is easy to simulate every MPC algorithm in the AMPC model. Namely, in the AMPC model, instead of sending a message to machine with id x , we can write a key-value pair keyed by x to the DHT. In the following round, each machine reads all key-value pairs keyed by its id. The limits on the number of queries and writes in the AMPC model are direct counterparts of the communication limits of the MPC model.

Due to known simulations of PRAM algorithms by MPC [26, 29], the AMPC model can also simulate existing PRAM algorithms from the EREW, CREW, CREW models [15], as well as stronger PRAM variants such as the MultiPrefix PRAM [39]. The simulations can be done using $O(1)$ rounds per PRAM step, and total space proportional to the number of processors.

It is also worth noting that a follow-up work [11] evaluated the AMPC algorithms empirically. The results show that the implementations of AMPC algorithms for various fundamental graph problems (including the MIS algorithm from this article) obtain significantly better running times, compared to the MPC implementations.

Minor changes. We note that the Adaptive MPC model as described here has two minor changes compared to the description in our conference version. The order of the complexity of all algorithms described in this article is unchanged by these modifications.

The first change is to model the system serving queries and responses as a distributed hash-table instead of a distributed data store. Note that we only use key-value semantics in our work, which guarantees that multiple queries for the same key always return the same result.

The second change is related to querying the values with the same key and gives a more precise modeling assumption that is beneficial for both clarity, and for studying lower-bounds in the model. Our previous model is the same as the “simplified model” above, without the restriction that the number of values associated with a given key is at most S . This change is a minor, but important technical modification to the model, which places an upper-bound on the total number of values associated with a given key. Below, we explain why this change is important.

Consider an algorithm that emits $O(N)$ values, all associated with the same key. To serve this data in the next round, the DHT must map each value to a set of consecutive distinct indices, likely by sorting. However, if the sorting was done in the MPC model with $O(S)$ space per machine, then it would provably require $\Theta(1/\epsilon)$ rounds [26, 40]. This factor of $O(1/\epsilon)$ in the round-complexity would not be accounted for in our previous model, but is captured here.³

As an example of how this property of our previous model could be used, consider the following algorithm for computing the XOR of N bits. In the first round, the algorithm maps over all bits, and if the bit is one, then it emits a value for the key 1. In the next round, a single machine performs a binary search over the indexed values for key 1, spending $O(\log N)$ queries to find the index of the last value corresponding to 1, which also reveals the parity of the number of ones, which is the XOR of the bit-string. This results in a 2-round algorithm for computing XOR. In our new model, XOR may be computed using $O(1/\epsilon)$ rounds.

It is easy to see that the second change only increases the round complexity by an $O(1/\epsilon)$ factor. In particular, we can simply transform each round of the previous model to $\Theta(1/\epsilon)$ rounds of the

³We are very grateful to Moses Charikar, Anna Ma, and Li-Yang Tan [18] for bringing this important detail in the model to our attention, which led to the change in our model description in the current version of this article.

new model as follows. For each set of values associated with the same key, we can sort the key-value pairs in $\Theta(1/\epsilon)$ rounds using the MPC algorithm of Goodrich et al. [26]. Then, we tag the keys with their index in the sorted array. This makes all the keys unique. Note that the increase in the round complexity, $\Theta(1/\epsilon)$, is a constant factor and hides in the O notation.

2.1 Realism

Let us now discuss the practical aspects of building a framework implementing the AMPC model.

Sequential queries. The AMPC model allows each machine to issue $O(S)$ queries within each round. Each query may depend on the results of the previous queries, which may cause the queries to be issued sequentially. It is natural to wonder how this aspect of the model may affect the parallelism. We believe that this assumption is realistic for the following reasons.

First, recent work on RDMA has shown that large-scale RDMA systems can be implemented efficiently, with many commercial systems, hardware and networking support for RDMA such as Infiniband [6], Omnipath [13, 20], and RoCE [7] being widely available. RDMA has been used in recent work on key-value storage systems [23, 28, 31, 34], database systems [49], and other related systems and studies [22, 37]. The latency for a remote memory access in recent RDMA systems is around 1–3 μ s under load, making it about 20 $\times$ more costly than a main-memory access [22].

Second, we observe that the per-machine latency can be further reduced by using the *parallel slackness* technique [44]. Assume there are $P = n^\epsilon$ machines, where $0 < \epsilon < 1$. Instead of running the algorithm directly on the machines, we split each machine into T^δ virtual machines of space $S = T^{\epsilon-\delta}$, for some small constant $0 < \delta < \epsilon$. Then, we execute the algorithm on the virtual machines. Each physical machine can simulate the virtual machines concurrently. In particular, when a virtual machine is stalled on a network access, the physical machine can swap it out and run a step of the next virtual machine, thereby hiding the latency of the query (note the purpose of hyper-threading in modern systems is also to perform latency-hiding).

Observe that this requires that the algorithm that is executed can be run on machines with space $S = T^{\epsilon-\delta}$. We remark that most of the MPC algorithms that require $S = T^{1-\Theta(1)}$ space per machine can run as long as $S = T^\epsilon$, for any constant ϵ (incurring only a constant factor increase in the number of rounds). In particular, this property holds for all algorithms presented in this article.

Fault tolerance. One of the key properties of the popular distributed computation frameworks implementing the MPC model is fault tolerance. This is achieved by storing the messages sent to each machine on persistent storage, usually a distributed file system [21, 32, 42, 48]. When a machine fails, it can simply restart its computation from scratch and read its input again. The AMPC model has similar characteristics. The key requirement is to implement each DHT in a fault-tolerant way [17]. After that, a failing machine can be simply replaced with a different machine that would perform the computation from scratch. Note that this crucially relies on the fact that the contents of $\mathcal{D}_{i-1}$ do not change within round i . Overall, we believe that the AMPC model can be implemented with similar fault-tolerance characteristics to the implementations of the MPC model.

Query contention. In the AMPC model each machine can issue $O(S)$ queries to the DHT in each round, but the machines are otherwise not restricted in accessing the DHT. In particular, all machines may be accessing the same key-value pair, which could potentially lead to contention. We observe that under the following natural assumptions, this should not be an issue.

- (1) $P = O(S^{1-\delta})$ for a constant $\delta > 0$. This is a realistic condition for at least two reasons. First, typically the number of machines P is up to few thousands, but the space S of each machine is billions of bytes of memory. Apart from that, in the popular distributed computation frameworks, there is a master machine that communicates with all other machines in each round [21, 32]. To make sure that the amount of this communication is bounded by $O(S)$, we must have $P = O(S)$.
- (2) The DHT is handled by P machines, each having $O(S)$ space. (In practice these can be either the machines doing the computation, or a different set of machines.) Each of these machines is capable of answering at most $O(S)$ queries, which is in line with our assumptions on the machines performing the computation.
- (3) The key-value pairs are randomly and independently assigned to the machines handling the DHT. The choice of the keys queried by each machine is independent of this mapping. This is a natural assumption, which basically says that the algorithm executed by a machine is independent of the technical details of the system.
- (4) Each worker machine queries for each key at most once. Note that this can be assumed without loss of generality, since the machines have sufficient space to cache the results of the queries they make.

The lemma below uses the above assumptions to analyze contention. It shows that each machine handling the DHT needs to answer $O(S)$ queries w.h.p. Note that in the statement of the lemma the balls are the key-value pairs and their weights give the number of times they are queried for. The proof is given in Section 10.

LEMMA 2.1. *Consider T balls, each having an integer weight between 0 and P . The total weight of all balls is T . Assume each of the balls is inserted uniformly at random to one of P bins. The bin for each ball is chosen independently of the choices for other balls and of the ball weight. Let $S = T/P$ and assume that $P = O(S^{1-\Omega(1)})$.*

Then, the total weight of balls in each bin is $O(S)$ with high probability.

Disallowing writes. In principle, the AMPC model could allow not only random reads, but also writes within a round. However, this would have a number of undesirable consequences. First of all, failure recovery becomes much more complex. When a machine fails, it cannot simply restart its computation within a round from scratch, since the contents of the DHT may have changed since the round has started. Moreover, if writes were allowed, then the machines could wait for the result of the computation of other machines, but the model would not account for this waiting time. In fact, one round of AMPC with writes could be used to simulate $O(S)$ steps of a PRAM algorithm.

3 PRELIMINARIES

We denote an unweighted graph by $G(V, E)$, where V is the set of vertices and E is the set of edges in the graph. The number of vertices in a graph is $n = |V|$, and the number of edges is $m = |E|$. Vertices are assumed to be indexed from 0 to $n - 1$. For undirected graphs, we use $N(v)$ to denote the neighbors of vertex v and $\deg(v)$ to denote its degree. We use D to refer to the diameter of the graph, or the longest shortest path distance between any vertex s and any vertex v reachable from s . We assume that there are no self-edges or duplicate edges in the graph.

Most of the algorithms presented in this article take a graph as an input. In this case, the input size N is equal to $n + m$, that is the total number of vertices and edges in the graph. We also assume that $S = n^\epsilon$, where ϵ is a constant, such that $\epsilon \in (0, 1)$. Depending on the algorithm, the total space of all machines required by the algorithm is either $\Theta(N)$ or $\Theta(N \log N)$. In the latter case,

the number of machines is $P = \Theta((N \log N)/n^\epsilon)$. In some cases, we may obtain a faster algorithm if $T = n^{1+\Omega(1)}$.

The round complexities of our algorithms have a polynomial dependency on $1/\epsilon$. However, since we assume ϵ to be a constant, for algorithms running in $\omega(1)$ rounds we hide this dependency in the $O(\cdot)$ notation.

We say that an event happens *with high probability* (w.h.p.) if it happens with probability at least $1 - 1/n^c$, where the constant c can be made arbitrarily large.

In the descriptions of our algorithms, we highlight the parts that rely on AMPC model features. The remaining parts can be implemented in MPC model using standard primitives, such as sorting, duplicate removal, and so on.

4 WARM-UP: THE 2-CYCLE PROBLEM

In this section, we give an algorithm that given an instance of the 2-CYCLE problem, solves it in $O(1)$ rounds w.h.p. Recall that an instance of the problem is either a single cycle on n vertices, or two cycles on $n/2$ vertices each. The problem is to detect whether the graph consists of one cycle, or two cycles. Our algorithm for the problem runs over a number of rounds. In each round, we sample each vertex with probability $n^{-\epsilon/2}$. Our goal is to replace the paths between sampled vertices with single edges, which can be viewed as contracting the original graph to the samples. We perform this contraction by traversing the cycle in both directions for each sampled vertex until we hit another vertex that is sampled. Traversing the cycle is implemented using the adaptivity of the model. To achieve this, it suffices to store in the DHT the graph, in which the sampled vertices are marked. Using a sample probability of $n^{-\epsilon/2}$, we can show that the number of vertices shrinks by a factor of $n^{\epsilon/2}$ in each round w.h.p. Therefore, after $O(1/\epsilon)$ rounds, the number of remaining vertices and edges is reduced to $O(n^\epsilon)$, at which point, we can fit the graph in the memory of a single machine and solve the remaining problem in a single round.

LEMMA 4.1. *Let G be a graph consisting of cycles, and let n be the initial number of vertices in G . Consider a cycle with size $k = \Omega(n^\epsilon)$ in some iteration of the loop of $\text{SHRINK}(G, \epsilon, O(1/\epsilon))$. The size of this cycle shrinks by at least a factor of $n^{\epsilon/2}$ after this iteration w.h.p.*

ALGORITHM 1: $\text{SHRINK}(G = (V, E), \delta, t)$

- (1) For $i := 1, \dots, t$:
 - (a) Sample each vertex independently with probability $n^{-\delta/2}$. Denote the set of sampled vertices by M . Randomly distribute the vertices of M to the machines.
 - (b) For each sampled vertex v traverse the cycle in each direction until a sampled vertex is reached. Let l_v and r_v be the sampled vertices, where the each of two traversals finishes.
- (2) Return the graph $(M, \{xl_x \mid x \in M\} \cup \{xr_x \mid x \in M\})$.

ALGORITHM 2: $\text{2-CYCLE}(G = (V, E))$

- (1) $G' := \text{SHRINK}(G, \epsilon, O(1/\epsilon))$
- (2) Solve the 2-CYCLE problem on G' , which has size $O(n^\epsilon)$ w.h.p. on a single machine.

PROOF. Consider the cycle, which by assumption has size $\Omega(n^\epsilon)$. Let the vertices in this cycle be labeled $1, \dots, k$, and let X_i be the random variable that is 1 if the i th vertex is marked, and 0

otherwise. Let $X = \sum_{i=1}^k X_i$. Since the sampling probability in each round is $n^{-\epsilon/2}$, we have that $E[X] = \frac{k}{n^{\epsilon/2}} = \Omega(\log n)$. Applying a Chernoff bound, we have the following:

$$P(X \geq (1+c)E[X]) \leq \exp(-c^2 \log n/3) = n^{-c^2/3}.$$

The number of vertices in this cycle in the next iteration is just the number of cycle vertices that are sampled. By the argument above, the number of sampled vertices is $O(E[X]) = O(\frac{k}{n^{\epsilon/2}})$ w.h.p. which is a shrink by a factor of $\Omega(n^{\epsilon/2})$. $\square$

LEMMA 4.2. *Let G be an n -vertex graph consisting of cycles and let $G' = \text{SHRINK}(G, \epsilon, O(1/\epsilon))$. Then G' is a graph that can be obtained from G by contracting edges. The length of each cycle in G' is $O(n^\epsilon)$.*

PROOF. From the pseudocode, it is clear that G' can be obtained from G by contracting edges. By Lemma 4.1, the size of a cycle with size $k = \Omega(n^\epsilon)$ shrinks by a factor of $n^{\epsilon/2}$ w.h.p. after an iteration of the loop in SHRINK. Hence, after the first iteration, the length of each cycle(s) is $O(n^{1-\epsilon/2})$ w.h.p, and after the i th iteration, the length of the cycle(s) is $O(n^{1-i\epsilon/2})$. After $\frac{2(1-\epsilon)}{\epsilon}$ iterations, the size of the cycle(s) is $O(n^\epsilon)$. $\square$

LEMMA 4.3. *The total communication of each machine is $O(n^\epsilon)$ in each round w.h.p., where n is the initial number of vertices in G .*

PROOF. Following the proof of Lemma 4.2, in the i th iteration, there are $O(n^{1-i(\epsilon/2)})$ vertices remaining, and $O(n^{1-(i+1)\epsilon/2})$ vertices sampled w.h.p. Recall that there are $P = \Omega(n^{1-\epsilon})$ machines, so each machine receives at most $O(n^{1-(i+1)(\epsilon/2)-(1-\epsilon)}) = O(n^{\epsilon/2})$ samples w.h.p.

Observe that the length of each traversal is a geometric random variable with $p = n^{-\epsilon/2}$. Let Z_i be the number of vertices that the i th sample traverses until it hits a sampled vertex. Let $Z = \sum_{i=0}^{n^{\epsilon/2}} Z_i$ be a random variable measuring the total amount of queries used by the samples assigned to some machine. Using a Chernoff bound for geometric random variables, and the fact that $E[Z] = n^\epsilon$, for some arbitrary large constant c we have

$$P(Z \geq (1+c)E[Z]) \leq \exp(-c^2 E[Z]/3) \leq \exp(-c^2 \log n) = O(n^{-c}).$$

Therefore, the total number of queries sent by the machine is $O(n^\epsilon)$ on any round w.h.p. $\square$

Combining the previous lemmas gives us the following theorem:

THEOREM 1. *There is an AMPC algorithm solving the 2-CYCLE problem in $O(1/\epsilon)$ rounds w.h.p. using $O(n^\epsilon)$ space per machine and $O(n)$ total space.*

5 MAXIMAL INDEPENDENT SET

In this section, we prove the following result.

THEOREM 2. *For any constant $\epsilon \in (0, 1)$, there exists an AMPC algorithm that finds a maximal independent set of a given graph $G(V, E)$ in $O_\epsilon(1)$ rounds. The algorithm is randomized and fails with arbitrarily small constant probability, uses $O(n^\epsilon)$ space per machine and $O(m)$ total space.*

Our algorithm for Theorem 2 finds the *lexicographically first MIS* over a random permutation of the vertices. Given a permutation $\pi : [n] \rightarrow [n]$, the lexicographically first MIS according to π , which we denote by $\text{LFMIS}(G, \pi)$, is constructed by processing the vertices in the order of π and adding a vertex v to the MIS if and only if it does not lead to a conflict. That is, v joins the MIS if none of its already processed neighbors have joined it before.

This “greedy”-type construction of $\text{LFMIS}(G, \pi)$ is not directly applicable in our model if we require a small number of rounds. Instead, we have to parallelize this algorithm. For this, we first

recall a query process in References [36, 47] to determine whether a given vertex v belongs to $\text{LFMIS}(G, \pi)$, then show how we can use this query process to efficiently construct $\text{LFMIS}(G, \pi)$ in $O(1)$ rounds.

For a vertex v , define $f(v, \pi)$ to be the indicator function for whether $v \in \text{LFMIS}(G, \pi)$. That is, $f(v, \pi) = 1$ if $v \in \text{LFMIS}(G, \pi)$ and $f(v, \pi) = 0$ otherwise. Note that we hide the dependence on G in the $f(v, \pi)$ notation for brevity. Determining the value of $f(v, \pi)$ can be done in the following recursive way.

ALGORITHM 3: $f(v, \pi)$

- (1) Let $u_1, \dots, u_d$ be the neighbors of v sorted such that we have $\pi(u_1) < \dots < \pi(u_d)$.
- (2) For $i = 1, \dots, d$:
 - If $\pi(u_i) < \pi(v)$:
 - if $f(u_i, \pi) = 1$, then return $f(v, \pi) = 0$.
- (3) Return $f(v, \pi) = 1$.

Now, for a vertex v , define $q_\pi(v)$ to be the number of (recursive) calls to the algorithm above for determining the value of $f(v, \pi)$. A result of Yoshida et al. [47] can be succinctly summarized as:

PROPOSITION 5.1. *If permutation π is chosen uniformly at random, then $\mathbb{E}_\pi[\sum_{v \in V} q_\pi(v)] \leq m + n$.*

Suppose that we fix a random permutation over the vertices, say, by each vertex v picking a random real $\rho_v \in [0, 1]$ and constructing π by sorting the vertices based on ρ . Having this, the natural way of implementing Algorithm 3 in the AMPC model is to assign each vertex v to a specific machine and simulate $f(v, \pi)$ there by exploring the relevant neighborhood of v according to the query process. The main challenge is that while Proposition 5.1 bounds the total communication by $O(m)$ in expectation, there is no upper bound on $q_\pi(v)$ for a given vertex v . Note that the machine responsible for vertex v has to query $q_\pi(v)$ vertices. Since the number of such queries within a round is bounded by the local space on the machines (S), this idea does not work if $q_v = \omega(n^\epsilon)$ for some vertex v . To ensure that we do not violate the per machine capacity on the number of queries and the total communication size, we use the iterative and truncated variant of the algorithm shown in Algorithm 4.

It is clear from the description that $\text{is-in-MIS}(v)$ either equals the value of $f(v, \pi)$ or is unknown and that once Algorithm 4 terminates it returns $\text{LFMIS}(G, \pi)$. It remains to analyze the number of required iterations.

ALGORITHM 4: $\text{MAXIMALINDEPENDENTSET}(G = (V, E))$

- (1) Fix a permutation $\pi : [n] \rightarrow [n]$ on the vertices uniformly at random.
- (2) Each vertex v is assigned to a machine μ_v chosen uniformly at random from m/n^ϵ machines.
- (3) For each vertex v , set $\text{is-in-MIS}(v) \leftarrow \text{unknown}$.
- (4) Repeat while the graph is not empty:
 - (a) For each vertex v remaining in the graph, μ_v runs $\text{TRUNCATEDQUERY}(v, \pi, n^\epsilon)$. If $\text{is-in-MIS}(v) = \text{true}$, then for every neighbor u of v , set $\text{is-in-MIS}(u) \leftarrow \text{false}$ and remove u from the graph as well.

ALGORITHM 5: TRUNCATEDQUERY(v, π, c)**Input:** a vertex v , a permutation π and a capacity c on the number of (recursive) queries allowed.**Output:** the number of conducted queries.

- (1) If $c = 0$, then return 0.
- (2) Let $u_1, \dots, u_d$ be the remaining neighbors of v sorted such that $\pi(u_1) < \dots < \pi(u_d)$.
- (3) $q \leftarrow 1$
- (4) For $i = 1, \dots, d$:
 - (a) If $\pi(v) < \pi(u_i)$, set is-in-MIS(v) = true and return q .
 - (b) $q \leftarrow q + \text{TRUNCATEDQUERY}(u_i, \pi, c - q)$
 - (c) If as a result of the query above, we get that u_i is in the MIS, then set is-in-MIS(v) = false and return q .
 - (d) If $q \geq c$, then return q .

As we will later discuss, each iteration of the Line 4 loop in Algorithm 4 can be implemented in $O(1)$ rounds of AMPC. Thus, it suffices to bound the number of iterations of this loop to bound the round complexity of the algorithm. In what follows, we simply say “iteration” to refer to iterations of this loop. We use the following lemma to bound the number of iterations by $O(1/\epsilon)$. For simplicity, we say a vertex v is *settled* if is-in-MIS(v) $\neq$ unknown.

LEMMA 5.2. *By the end of iteration i , any vertex v with $q_\pi(v) \leq n^{i\epsilon/2}$ will be settled.*

We prove Lemma 5.2 by induction. For the base case with $i = 1$, it is clear that by the end of iteration 1, any vertex with $q_\pi(v) \leq n^{\epsilon/2}$ will be settled, since the capacity on the number of queries is n^ϵ . Suppose that this holds for any iteration $i \leq r - 1$. We show that by the end of iteration r , any vertex v with $q_v \leq n^{r\epsilon/2}$ will be settled.

For any vertex v , define $s_i(v) = \langle w_1, w_2, \dots, w_q \rangle$ to be the sequence of vertices for which function TRUNCATEDQUERY is (recursively) called within iteration i sorted in the same order that they were visited. Observe that for a vertex v that is not settled by the end of iteration i , $s_i(v)$ has to have exactly n^ϵ vertices, since this is the cap that we set for any vertex in Algorithm 4. The following claim will be useful in proving Lemma 5.2.

CLAIM 5.3. *Fix an arbitrary vertex v and let D be an upper bound on the depth of recursions within TRUNCATEDQUERY(v, π, c). If v is not settled by the end of iteration r , then the following holds if we denote $s_r(v) = \langle w_1, \dots, w_q \rangle$: For any integer $j \in [q - D]$, there is at least one vertex $w \in \{w_j, w_{j+1}, \dots, w_{j+D}\}$ such that $q_\pi(w) \geq n^{(r-1)\epsilon/2}$ and that w is settled during iteration r without calling TRUNCATEDQUERY on any other vertex.*

PROOF. Since D is an upper bound on the depth of recursions, there is at least one vertex $w_{k+1} \in \{w_{j+1}, \dots, w_{j+D}\}$ that is not queried by w_k (otherwise the query depth from w_j to w_{j+D} would be $D + 1$, which is a contradiction). We show that vertex w_k satisfies the requirements.

Observe that in Algorithm 5, we call TRUNCATEDQUERY only on vertices that are not already settled. Since $w_k \in s_r(v)$, vertex w_k should not have been settled by the end of iteration $r - 1$. By the induction hypothesis, this means that $q_\pi(w_k) > n^{(r-1)\epsilon/2}$. Furthermore, the fact that vertex

w_{k+1} , the successor of w_k in $s_r(v)$, is not queried by w_k , means that w_k has not queried any vertices and thus it should be settled within round r , proving the claim. $\square$

Now, take vertex v with $q_v \leq n^{r\epsilon/2}$ and suppose for contradiction that v is not settled by round r . Recall that this means $|s_r(v)| = \Theta(n^\epsilon)$. By Claim 5.3, there are at least $\Omega(|s_r(v)|/D)$ vertices u with $q_u \geq n^{(r-1)\epsilon/2}$ that are settled within round r without querying any vertex. This saves v from querying $\Omega(\frac{|s_r(v)|}{D} \cdot n^{(r-1)\epsilon/2}) = \Omega(\frac{n^\epsilon}{D} \cdot n^{(r-1)\epsilon/2})$ vertices compared to the original query process $f(v, \pi)$. The recursion depth can be bounded in a straightforward way by adapting the arguments in Reference [14] (see Claim 5.6 below) by $D = O(\log^2 n)$. Thus, a total of $\Omega(\frac{n^\epsilon}{\log^2 n} \cdot n^{(r-1)\epsilon/2}) \gg n^{r\epsilon/2}$ queries should not be sufficient to settle the status of v . This is a contradiction, since we assumed $q_v \leq n^{r\epsilon/2}$.

To conclude this section, it thus only remains to show how the recursion depth can be upper bounded by $O(\log^2 n)$. The argument for this is similar to the analysis of Reference [14] for bounding the round complexity of a parallel implementation of the randomized greedy MIS algorithm, though we note that applying their result as a black-box does not bound the depth of recursions. We use the following lemmas proved in Reference [14]:

LEMMA 5.4 ([14, COROLLARY 3.2]). *Let Δ denote the maximum degree of G . Suppose that the ordering on vertices is uniformly random, and for any integer $i \in [\lceil \log \Delta \rceil]$, consider the $(2^i \log(n)/\Delta)$ -fraction of vertices with the highest priorities. If the greedy MIS on these vertices along with all of the vertices adjacent to this MIS are removed from G , then the remaining graph will have maximum degree $\Delta/2^i$ with high probability.*

LEMMA 5.5 ([14, COROLLARY 3.4]). *Suppose that all vertices in a graph have degree at most Δ , and consider a random $O(\log(n)/\Delta)$ -fraction of the vertices. Suppose also that the vertices in this subset are randomly ordered. Then the longest monotone path (with regards to the ordering) in the induced graph, has length $O(\log n)$ with high probability.*

We are now ready to formally prove our claim that the recursion depths are polylogarithmic.

CLAIM 5.6. *With high probability over the randomization of permutation π , for any vertex v , the recursion depth in determining $f(\pi, v)$ via Algorithm 3 is bounded by $O(\log^2 n)$.*

PROOF. For any $i \in [k]$ with $k = \lceil \log_2 \Delta \rceil$, let F_i be the first $2^i \log(n)/\Delta$ fraction of vertices with the highest priorities in G and note that $F_1 \subset F_2 \subset \dots \subset F_k = V$. Let $V_0 = V$ and for $i \geq 1$, let V_i be the subset of vertices in G remained after removing the MIS nodes in F_i and all vertices adjacent to them (thus $V = V_0 \supseteq V_1 \supseteq \dots \supseteq V_k = \emptyset$). From Lemma 5.4, the induced graph $G[V_i]$ has maximum degree $\Delta/2^i$ w.h.p. Applying Lemma 5.5 on subsample $F_{i+1} \cap V_i$ of graph $G[V_i]$, yields that the maximum monotone path length in $G[F_{i+1} \cap V_i]$ is $O(\log n)$, w.h.p.

Now suppose for contradiction that the recursion depth of a vertex v is $l = c \log^2 n$ for a sufficiently large c . Thus, there is a recursion path $\langle v_0 = v, v_1, v_2, \dots, v_l \rangle$ such that $\pi(v_i) \leq \pi(v_{i-1})$ for any i . For each vertex u in the path, let $\ell(u)$ be the maximum i such that $u \in F_i$. Note that for any i , $\ell(v_i) \leq \ell(v_{i-1})$, since $\pi(v_i) \leq \pi(v_{i-1})$. Since there are at most $k = O(\log \Delta) = O(\log n)$ choices for $\ell(v_i)$, and that the length l of this recursion path is larger than $c \log^2 n$, we get that there are at least $s = c' \log n$ consecutive vertices $v_j, \dots, v_{j+s-1}$ in the path, such that $\ell(v_j) = \dots = \ell(v_{j+s-1}) = \ell$. We claim that there must be a vertex $u \in \{v_j, \dots, v_{j+s-2}\}$ (note that this set deliberately excludes the last vertex v_{j+s-1} of the subpath) such that $u \notin V_{\ell-1}$, i.e., u must have a neighbor in $F_{\ell-1}$ that joins the MIS. To see this, observe that otherwise any vertex u in set $u \in \{v_j, \dots, v_{j+s-2}\}$ satisfies $u \in F_\ell \cap V_{\ell-1}$, contradicting that the maximum monotone path in $G[F_\ell \cap V_{\ell-1}]$ is $O(\log n)$, since c'

can be made arbitrarily large. Finally, the fact that u has a neighbor in $F_{\ell-1}$ that joins the MIS, combined with the query process of Algorithm 3 that greedily queries the vertices of higher priority first, implies that u will not query another vertex w with $\ell(w) = \ell$, contradicting that the recursion path continues any further from u . As a result, our assumption that there can be a recursion depth of size $c \log^2 n$ for any large constant c cannot hold and the maximum recursion depth is indeed $O(\log^2 n)$. $\square$

6 UNDIRECTED GRAPH CONNECTIVITY

In this section, we provide an implementation of the recent undirected connectivity algorithm due to Andoni et al. [2] in the AMPC model in $O(\log \log_{T/n} n)$ rounds. This improves upon the original result by a factor of $\log D$, where D is the diameter of the input graph. Note that by our assumption $T = O(n+m)$. However, if we had $T = n^{1+\Omega(1)}$, then our algorithm would run in a constant number of rounds.

At a high level, we show that using the AMPC model, we can speed-up the central procedure in the algorithm of Andoni et al., so that it runs in only $O(1)$ rounds, compared to $O(\log D)$ rounds in the original algorithm. However, for completeness we provide the entire algorithm.

The Andoni et al. Algorithm. If the total amount of space used is not a concern, then undirected connectivity can easily be solved in $O(\log D)$ rounds in the MPC model with sublinear space per machine as follows. Let G^2 be a *squared* graph, that is a graph obtained by adding edges between each vertex and all its 2-hop neighbors. By repeating the squaring $\lceil \log_2 D \rceil$ times, we obtain a graph, where each connected component is transformed into a clique. The issue with the squaring approach is that the total amount of space used to store the resulting graph can be much as much as $\Theta(n^2)$, which is prohibitive for sparse graphs.

The main idea behind the Andoni et al. algorithm is to implement the graph exponentiation idea in a more space-efficient manner. The algorithm runs over a number of phases and in each phase contracts the graph, reducing the number of vertices. The algorithm also assigns each vertex a *budget* d , making sure that the sum of budgets of all vertices is $O(T)$. As the number of vertices decreases, the budgets are increased. The budgets of all vertices are equal to each other at all times.

Let us now describe one phase. The first step of a phase increases the degree of each vertex to at least d , by adding edges within each connected component (and without adding parallel edges). After that every vertex is sampled as a *leader* independently with probability $\Theta(\log n/d)$ (assume $d = \Omega(\log n)$ for simplicity). The non-leader vertices are then contracted to leaders (a leader exists in each non-leader's neighborhood w.h.p.). The total number of vertices after contraction is $\tilde{O}(n/d)$. Therefore, in the next round, the budgets can be increased to $d^{1.4}$ and the sum of budgets will be $O(nd^{0.4}) = O(T)$. As the budgets increase double exponentially, $O(\log \log n)$ phases are required before each connected component shrinks to one vertex.

In the Andoni et al. algorithm, each phase is implemented by using $O(\log D)$ rounds of squaring. In each round, each vertex that has degree less than its budget connects itself to its 2-hop (or a sufficiently sized subset of its 2-hop if the 2-hop is too large). In the worst case, the algorithm requires $O(\log D)$ rounds to make the degrees equal to the budget.

Implementing Andoni et al. in AMPC. We show how to implement each phase in $O(1)$ rounds using the adaptivity in the AMPC model. Recall that the goal of a phase is to increase the degree of each vertex in the current graph to d , the current per-vertex budget. In the AMPC model, this can be achieved by exploring the local neighborhood of each vertex in one round (e.g., using BFS). The search stops once it visits d vertices.

One possible worry is that there may be significant overlap amongst the neighbors that we query while probing from a vertex. Namely, running BFS until it visits d vertices may require visiting $\Omega(d^2)$ edges. We cope with this issue by setting the budget to be the square root of the space available per each vertex. Therefore, when a vertex attempts to increase its degree to d in a phase, we can afford to issue d^2 queries.

Once d reaches $n^{\epsilon/2}$, we no longer increase it, which assures that the number of queries issued per vertex is at most $O(n^\epsilon)$. From this point the number of vertices decreases by a factor of $d = \Omega(n^{\epsilon/2})$ in each iteration. It follows that after $O(1/\epsilon)$ iterations each connected component will be contracted to a single vertex. The AMPC algorithm for a phase is given in Algorithm 6.

ALGORITHM 6: INCREASEDEGREES($G = (V, E), d$)

- (1) For each vertex v , run a BFS from v , until it visits d vertices or visits the entire connected component of v . Let N_v be the set of visited vertices.
- (2) Return $(V, E \cup \{vx \mid v \in V, x \in N_v\})$

Algorithm 6 can be naturally implemented in the AMPC model. First, the graph is written to the DHT. Then, the vertices are randomly assigned to machines, and each machine runs the BFS for its assigned vertices. We now analyze this implementation.

LEMMA 6.1. *Algorithm 6 runs in $O(1)$ rounds of AMPC. The total number of queries is $O(nd^2)$, where n denotes the number of vertices of G . If the number of machines is $n^{1-\epsilon}d^2$, then the maximum number of queries a machine makes is $O(n^\epsilon)$ w.h.p.*

PROOF. The algorithm clearly runs in $O(1)$ rounds. To bound the number of queries, note that in the worst case we make a query for every edge between a pair of vertices that are visited by the BFS. Since we stop after visiting d vertices, the maximum number of queries we make is $O(d^2)$ and therefore the total number of queries is $O(n^\epsilon d^2)$.

Next, we use a Chernoff bound to bound the maximum number of queries a machine makes. Fix a machine i and let binary random variable X_v indicate whether vertex v is assigned to machine i . Let $X = \sum_{v \in V} X_v$. Note that X is the total number vertices assigned to machine i , and $E[X] = \frac{n^\epsilon}{d^2} \geq n^{\epsilon/3}$. By a Chernoff bound, we have

$$P(X \geq 2E[X]) \leq \exp\left(-\frac{E[X]}{3}\right) = \exp(-n^{\epsilon/3}/3).$$

A union bound over all $d^2 n^{1-\epsilon}$ machines implies that the total number of queries of each machine is upper bounded by $E[X] \cdot d^2 = O(n^\epsilon)$, with high probability. $\square$

By plugging in the above procedure to the algorithm of Andoni et al., we are able to reduce the running time by a factor of D .

In addition, we also show how to ensure that the algorithm runs correctly with high probability, as opposed to the constant probability of success of the original algorithm. To achieve this, we need to assure that after each vertex is chosen as a leader with probability $\Theta(\log n/d)$, every vertex has a leader among its neighbors with high probability.

To this end, we need to assure that the degree of each vertex in the graph before sampling is $\Omega(\log n)$. We can achieve this by running Algorithm 6 with $d = \Omega(\log n)$. By Lemma 6.1, the number of queries issues by the algorithm is nd^2 , which does not exceed the allowed number of

queries $T = O(m)$, whenever $m = \Omega(n \log^2 n)$. Hence, we handle the case when $m = o(n \log^2 n)$ separately.

When $m = o(n \log^2 n)$, we use the algorithm in Reference [10] to shrink the number of vertices by a factor of $\Omega(\log^2 n)$ in $O(\log \log_{T/n} n)$ rounds w.h.p. This procedure is implementable in the MPC model and hence implementable in the AMPC model as well. The following lemma summarizes the properties of the algorithm.

LEMMA 6.2 ([10]). *There exists an $O(\log \log n)$ round MPC algorithm using $O(n^\epsilon)$ space per machine and $O(m)$ total space that with high probability, converts any graph $G(V, E)$ with n vertices and m edges to a graph $G'(V', E')$ and outputs a function $f : V \rightarrow V'$ such that:*

- (1) *The number of non-isolated vertices in G' is $O(n/\log^2 n)$.*
- (2) *$|E'| \leq |E|$.*
- (3) *For any two vertices u and v in V , vertices $f(u)$ and $f(v)$ in V' are in the same component of G' if and only if u and v are in the same component of G .*

Undirected Connectivity. We assume that $m = \Omega(n \log^2 n)$ in what follows. Next, we describe the undirected connectivity algorithm in the AMPC model.

ALGORITHM 7: CONNECTIVITY($G=(V, E)$)

- (1) Let $d = \sqrt{T/n}$. Let $M : V \rightarrow \mathbb{Z}$ be a mapping from vertices to a unique identifier for their component (initially mapping each vertex to its own id).
- (2) While G contains at least one edge, perform the following steps:
 - (a) Randomly assign the vertices in G to machine and call INCREASEDEGREES(G, d) for each vertex. Add the edges found to G .
 - (b) Sample each vertex with probability $\Theta(\log n/d)$ to be a leader.
 - (c) Let G' to be the graph formed by contracting each vertex v to a leader in its neighborhood. If its degree is less than d , then contract it to the neighbor with lowest id. Let M' be a mapping $V(G) \rightarrow V(G')$ mapping each vertex to the vertex it contracts to in G' . Finally, update G to be G' .
 - (d) Update $d = \min(d^{1.4}, n^{\epsilon/3})$.
 - (e) Update M by mapping each $v \in V$ to $M'(M(v))$.
- (3) Output M

THEOREM 3. *Algorithm 7 computes the connected components of an undirected graph in $O(\log \log_{T/n} n + 1/\epsilon)$ rounds of AMPC w.h.p. where the total space $T = \Omega(m + n)$.*

PROOF. For simplicity let us assume that the graph is connected. Otherwise, we can apply the analysis to each connected component. First, we show that the sum of all *squared* budgets is $O(T)$ at every step of the algorithm. Indeed, this is true in the beginning, and in each iteration the number of nodes decreases by a factor of $d/\log n$ w.h.p. while the per-vertex squared budgets are increased by a factor of $d^{0.8}$. Clearly, the algorithm also assures that $d = O(n^{\epsilon/2})$. Since we assumed that $m = \Omega(n \log^2 n)$, in the first iteration we have $d = \Omega(\log n)$. This implies that in this iteration, as well as all following ones, calling INCREASEDEGREES will increase degrees to at least d , except for the vertices contained in connected components of size at most d . Hence, w.h.p. each vertex

either has a neighbor who is a leader, or belongs to a connected component of size at most d , which INCREASEDEGREES turns into a clique.

From the analysis in Reference [2] it follows that the number of iterations of the algorithm until $d = \Omega(n^{\epsilon/3})$ is $O(\log \log_{T/n} n)$. Once d reaches $\Omega(n^{\epsilon/3})$, the number of vertices in each connected component shrinks by a factor of $\Omega(n^{\epsilon/3})$ in each iteration. Hence, the total number of iterations is $O(\log \log_{T/n} n + 1/\epsilon)$. Each iteration can be implemented in $O(1)$ rounds (see Lemma 6.1).

It remains to analyze the number of queries in each round. Clearly, it suffices to analyze the number of queries for the part of the algorithm that cannot be implemented in the MPC model, which is Algorithm 6. By Lemma 6.1 and the fact that the sum of squared budgets is $O(T)$, we get that the total number of queries in one round is $O(T) = O(m)$. Moreover, the maximum number of queries per machine is $O(n^\epsilon)$ w.h.p. $\square$

7 MINIMUM SPANNING FOREST

In this section, we describe a **minimum spanning forest (MSF)** algorithm with a round-complexity of $O(\log \log_{T/n} n)$ w.h.p. in the AMPC model. The input to the algorithm consists of an undirected weighted graph, $G = (V, E)$. For simplicity, we assume that all edge weights are distinct to ensure that there is a unique MSF (the assumption can easily be met by, for example, breaking ties by the ids of the endpoints). The algorithm is similar in spirit to the connectivity algorithm in Section 6. In the first phase of the algorithm, each vertex computes a subset of the minimum spanning forest using Prim's algorithm [19]. In the second phase, we randomly sample leaders like in the connectivity algorithm, and contract each non-leader vertex to a leader in its neighborhood.

When $m = o(n \log^2 n)$ we use the algorithm of Lemma 6.2, which shrinks the number of vertices by a factor of $\Omega(\log^2 n)$ in $O(\log \log_{T/n} n)$ rounds w.h.p. The procedure executes a number of steps, each of which shrinks the number of vertices by a constant factor w.h.p. We invoke each step of the procedure with the graph formed by each vertex and the lowest weight edge incident to the vertex (like in Borůvka's algorithm) so that the contractions performed will be along edges in the MSF. Since the procedure decreases the number of vertices by a factor of $\Omega(\log^2 n)$ in $O(\log \log_{T/n} n)$ rounds w.h.p., in what follows we assume that $m = \Omega(n \log^2 n)$.

Algorithm 8 gives pseudocode for the degree increasing procedure in our minimum spanning forest algorithm. It takes as input a graph G , and a degree-bound d , and runs Prim's algorithm from each vertex, v , until the size of the MSF rooted at v is d .

ALGORITHM 8: MSFINCREASEDEGREE($G = (V, E), d$)

- (1) For each vertex, v , assigned to a machine, compute its local spanning forest, F_v , and $E(v)$ as follows:
 - (a) Initially, $F_v = \{v\}, E(v) = \{\}$
 - (b) While $|F_v| < d$, select the minimum weight edge, (v_1, v_2) s.t. $v_1 \in F_v$ and $v_2 \notin F_v$, and add v_2 to F_v . Add (v_1, v_2) to $E(v)$. If such an edge does not exist, halt, since F_v contains v 's entire component.
- (2) Output F_v and $E(v)$ for each vertex v .

ALGORITHM 9: MINIMUMSPANNINGFOREST($G = (V, E)$)

- (1) Let $E_{MSF} = \emptyset$ and $M = \{w \rightarrow e \mid e = (u, v, w) \in E\}$ be a mapping from edge weights to edges.
- (2) Let $G_C = G, d = \sqrt{T/n}$
- (3) While G_C is non-empty, perform the following steps:
 - (a) Randomly assign the vertices in G_C to machines, and call MSFINCREASEDEGREE(G_C, d) for each vertex. The output is a set of neighbors, F_v and a set of minimum spanning forest edges, $E(v)$ for each $v \in V$.
 - (b) $E_{MSF} = E_{MSF} \cup \{M[E(v)] \mid v \in V\}$. Note that we apply M to the MSF edges found in this round in order to add the original edge in G corresponding to some edge G_C .
 - (c) Sample each vertex with probability $\Theta(\log n/d)$ to be a leader.
 - (d) Update G_C to the graph formed by contracting each vertex v to a leader in F_v . If its degree is less than d , then contract it to the neighbor in F_v with lowest id. Update M on the new graph.
 - (e) Update $d = \min(d^{1.4}, n^{\epsilon/3})$.
- (4) Output E_{MSF} .

Algorithm 9 is the main MSF routine. Recall that $m = \Omega(n \log^2 n)$ by assumption. We first call Algorithm 8 starting with an initial setting of $d = \Omega(\log n)$. We then perform leader selection to reduce the number of vertices in the graph by a factor of d . The decrease in the number of vertices allows us to increase d to $d^{1.4}$ in the next phase. Therefore, after $O(\log \log_{T/n} n)$ such phases, d becomes $O(n^{\epsilon/3})$, at which point, the algorithm terminates after $O(1/\epsilon)$ iterations.

LEMMA 7.1. *Algorithm 9 correctly outputs the minimum spanning forest of G .*

PROOF. The proof is by induction on the number of iterations, basing on the last iteration. In the base case, note that the iteration finds all remaining MSF edges, since G_C becomes empty after the final iteration by assumption. Since MSF edges are discovered using Prim's algorithm, and the edge weights are distinct, this iteration correctly computes the MSF of the graph.

In the inductive case, we just need to show that we find a valid subset of the MSF edges in this iteration, and that any MSF edge not yet found is still present in G_C . Again, since we discover edges using Prim's algorithm and due to each edge having a unique edge weight, we commit a valid subset of the MSF edges in this iteration. To show that any MSF edge not yet found is still preserved in G_C , observe that an edge (u, v) is only removed from G_C if both u and v merge to the same leader.

Suppose that an MSF edge (u, v) that is not committed is removed, i.e., is not represented in G_C . Then by the argument above, (u, v) merge to the same leader, and there is therefore a path P composed of MSF edges committed in this round going between $u, l(u)$ and v . Since we assumed that (u, v) is an MSF edge, some edge in P must have larger weight than (u, v) , but this is a contradiction, as it implies that the local searches using Prim's algorithm are incorrect. Thus, any removed edge cannot be an MSF edge, and so all MSF edges that are not found belong to G_C . $\square$

THEOREM 4. *Algorithm 9 computes the minimum spanning forest of an undirected graph in $O(\log \log_{T/n})$ rounds of AMPC w.h.p. where the total space $T = \Omega(m + n)$.*

PROOF. The proof follows by the same argument as the proof of Theorem 3. $\square$

Since we can compute a spanning forest of an undirected graph by assigning arbitrary distinct weights to the edges, we have the following corollary:

COROLLARY 7.2. *A spanning forest of an undirected graph can be found in $O(\log \log_{T/n})$ rounds of AMPC w.h.p. where the total space $T = \Omega(m + n)$.*

8 FOREST CONNECTIVITY

In this section, we present algorithms for solving *forest connectivity*, i.e., the undirected graph connectivity problem when the input is promised to be a forest. Our forest connectivity algorithm is based on the classic technique of transforming each forest into a cycle, via an Eulerian tour. After representing the forest as a set of Eulerian tours, finding the connectivity of the forest reduces to solving connectivity on a collection of cycles. Observing that the PRAM construction by Tarjan et al. [43] to transform a forest into a set of Euler tours (a collection of cycles) can be implemented in $O(1/\epsilon)$ rounds of AMPC leads us to the cycle connectivity problem.

Cycle connectivity. The *cycle connectivity problem* is to compute the connected components of a graph $G = (V, E)$ that contains a set of disjoint cycles. Our cycle connectivity algorithm, Algorithm 10, works as follows. We call $O(1/\epsilon)$ iterations of SHRINK (the same algorithm used in the 2-CYCLE problem) that reduces the size of the largest cycle to $O(n^\epsilon)$ with high probability. Then, we switch to a different algorithm that computes the connected components of the remaining cycles in one round using a total of $O(n \log n)$ communication with high probability. The algorithm first fixes a permutation π . Then, each vertex in the graph searches along one direction of the cycle until it either returns to itself or hits a vertex that appears before it in the permutation (i.e., has higher rank in π).

ALGORITHM 10: CYCLECONNECTIVITY($G = (V, E)$)

- (1) Let $G' := \text{SHRINK}(G, \epsilon, O(1/\epsilon))$
- (2) Fix a random permutation, π , over the vertices of G' and randomly assign the vertices to the machines.
- (3) For each vertex u assigned to a machine, search one direction of the cycle until either u fully traverses the cycle or u encounters a vertex v s.t. $\pi(v) < \pi(u)$. The vertex with lowest rank in π in a cycle is the representative for the cycle.

We obtain the following corollary by applying Lemma 4.1.

COROLLARY 8.1. *After $O(1/\epsilon)$ iterations of SHRINK with $\delta = \epsilon/2$, the largest cycle in the graph has size $O(n^{\epsilon/2})$ with high probability.*

LEMMA 8.2. *Consider a cycle of length $k = \Omega(\log n)$. The number of queries made by a vertex until it hits a vertex with higher rank in π is $O(\log k)$ in expectation.*

PROOF. Given an arbitrary vertex v , the probability that we have to query i vertices before finding a vertex is equal to the probability that the $i-1$ vertices following v in the cycle have lower rank than v in the permutation, and that the i th vertex has higher rank than v in the permutation. This probability is $\frac{1}{i} \cdot \frac{1}{i+1}$. The expected number of queries is therefore

$$\sum_{i=1}^k i \cdot \frac{1}{i} \cdot \frac{1}{i+1} = \sum_{i=1}^k \frac{1}{i+1} = H_k - 1 = O(\log k),$$

completing the proof. □

Next, we show that the total number of queries for large enough cycles is concentrated.

LEMMA 8.3. *In Algorithm 10, the total number of queries of a cycle of length $k = \Omega(\log n)$ is $O(k \log k)$ with high probability.*

PROOF. We observe that the amount of work done by the queries is identical to the work done by pivots selected in the randomized quicksort algorithm. To make the analogy concrete, consider first picking the highest priority element in π , and charge it k work for the queries it performs. Now, consider the second-highest priority element, which takes one of the remaining $k - 1$ slots. This element may perform $k - 1$ queries in the worst case, but it breaks the remaining array into two independent subproblems that cannot query each other. The query complexity follows from the analysis of randomized quicksort (see for example Reference [35]). $\square$

Next we bound the total number of queries of a machine. Since our bound on the total query complexity $O(n \log n)$ w.h.p., we use $O(n \log n)$ total space to perform the extra queries. Therefore, in what follows, we assume that we have $O(n^{1-\epsilon} \log n)$ machines. The proof boils down to analyzing a weighted version of the balls and bins problem, when we have an upper-bound on the maximum weight of any ball. Although a very similar problem is considered in Reference [41], it does not provide a high probability bounds on the maximum load of a bin.

LEMMA 8.4. *The total amount of queries performed by a machine in Algorithm 10 is $O(n^\epsilon)$ w.h.p.*

PROOF. The number of queries performed by SHRINK is at most $O(n^\epsilon)$ per round by Lemma 4.3. We now show that the total queries performed per machine in step (3) when traversing the cycles is $O(n^\epsilon)$ w.h.p.

To see this, we apply a lemma concerning the mean of a sample of size k from a universe of size n without replacement. Given a set of real numbers $c_1, \dots, c_n$ and the sample $X_1, \dots, X_k$, define $X = \sum_{i=1}^k X_i$, $\bar{c}_n = \frac{1}{n} \sum_{i=1}^n c_i$, $\sigma_n^2 = \frac{1}{n} \sum_{i=1}^n (c_i - \bar{c}_n)^2$, and, last, $\Delta_n = \max_{i=1}^n c_i - \min_{i=1}^n c_i$. The Hoeffding-Bernstein bound (see Lemma 2.14.19 from Reference [45]) states that

$$P(|X - \bar{c}_n| > t) \leq 2 \exp \left(- \frac{kt^2}{2\sigma_n^2 + t\Delta_n} \right).$$

Taking c_i to be the query complexity of the i th vertex, we now compute the terms defined above. Notice that by Lemma 8.3 we have $\sum_{i=1}^n c_i = O(n \log n)$ w.h.p., and so $\bar{c}_n = O(\log n)$, with high probability. Furthermore, by Corollary 8.1 we have $c_i < n^{\epsilon/2}$, with high probability. Therefore, we have $\forall i, (c_i - \bar{c}_n)^2 \leq n^\epsilon$, and $\sigma_n^2 \leq n^\epsilon$. Last, we have $\Delta_n \leq n^{\epsilon/2}$. Letting $t = c \cdot n^\epsilon$, we have

$$\begin{aligned} P(|X - \bar{c}_n| > cn^\epsilon) &\leq 2 \exp \left(- \frac{c^2 n^{3\epsilon}}{2n^\epsilon + cn^{3\epsilon/2}} \right) \\ &\leq 2 \exp(-cn^\epsilon) \\ &\leq 2n^{-c}, \end{aligned}$$

completing the proof. $\square$

Putting the lemmas above together, we have the following theorem:

THEOREM 5. *There exists an AMPC algorithm that solves the forest connectivity problem in $O(1/\epsilon)$ rounds of computation w.h.p. using $T = O(n \log n)$ total space w.h.p.*

8.1 Rooting Trees and List Ranking

We now consider the problem of computing a rooted tree given a set of edges representing a tree. We also present algorithms for computing tree properties, such as subtree sizes, preorder numbering, and subtree minima/maxima in this section. These algorithms are used in our 2-edge connectivity algorithm.

We start with the tree rooting problem. The input is a set of tree edges, and a vertex r to root the tree at. The output is a mapping $p : V \rightarrow V$ from vertices in the tree to their *parents* where $p(r) = r$. A classic parallel algorithm for rooting a tree is based on the Euler tour technique [43], which works as follows. First compute an Euler tour of the tree. Rooting the tree reduces to the *list ranking* problem, a classic problem in the parallel algorithms literature. We now show that a simple extension of our algorithm for the 2-CYCLE problem can be used to perform list ranking in the AMPC model.

The algorithm works as follows. It first computes the Euler tour of the tree, and then breaks it at an edge incident to the root, r , which unravels the cycle into a list. It then uses list ranking to compute the distance between each edge and the root of the list. Since each edge appears twice in the Euler tour, it also maintain cross-pointers between these edges.

ALGORITHM 11: LISTRANKING($L = \{v_0, \dots, v_N\}$)

- (1) Let $n = N$ and $r = 1$. Let $w_r(v)$ be the weight of vertex v on the r th round. Initially, $w_1(v) = 1$.
 - (2) While $n = \Omega(N^\epsilon)$, run the following steps.
 - (a) Sample each list element with probability $N^{-\delta/2}$ uniformly at random, and let S be the set of samples, including v_0 . Randomly distribute the samples to the machines.
 - (b) For each sampled vertex, v , traverse the list until another sampled vertex is hit and let $l(v)$ be number of vertices traversed before hitting the next sample. Update $w_r(v)$ to be $w_{r-1}(v) + l(v)$.
 - (c) Contract the graph to the samples, and update $r = r + 1$.
 - (3) Solve the remaining weighted list-ranking problem on a single machine using w_r . The output is the correct rank of all vertices that survived all r sampling rounds. Let $d_r(v) = w_r(v)$ for vertices that are active on the r th round and 0 otherwise.
 - (4) For $r' = r - 1 \dots 0$, run the following steps:
 - (a) Let v be a vertex sampled on round r in Step (a). Traverse the list until the next sampled vertex is hit, and update $d_{r'}(u_i) = d_{r'+1}(v) + \text{dist}(v, u_i)$ for all vertices $v, u_1, \dots, u_{l(v)}$.
 - (5) Output d_0 .
-

LEMMA 8.5. *Algorithm 11 correctly solves the list ranking problem.*

PROOF. First note that the contraction process of Algorithm 11 does not change the order of the vertices. Moreover, for each vertex v , each vertex u that is between v and the root may only contract to another vertex in this range. Thus, if after a round a vertex v survives, then the total weight of the vertices between v and the root remains the same. Next we use this observation to prove the statement of the lemma.

We prove this by an induction on the number of rounds, basing on the last round. For the base case, note that in after all contractions, we explicitly calculate the total weight of the vertices between each survived vertex and the root. As we mentioned before this is exactly the index of the survived vertex in the initial list. Next, assume that all the indices of the vertices in round $r' + 1$ are correctly set. Note that these are the sampled vertices in round r' . Hence, to find the index of a (not sampled) vertex v in round r' we just need to sum up the the index of its next sampled vertex with the total weight of the vertices between v and the sample vertex, as mentioned in the algorithm. This proves the induction step and completes the proof of the lemma. $\square$

THEOREM 6. *List ranking a list of length N can be performed in $O(1/\epsilon)$ rounds of AMPC w.h.p. Furthermore, the algorithm uses $O(N)$ total space w.h.p.*

PROOF. A proof identical to that of Lemma 4.1 shows that a list with size $k = \Omega(N^\epsilon)$ shrinks by a factor of at least $N^{\epsilon/2}$ after one iteration of the loop in Step 2. Therefore, $O(1/\epsilon)$ iterations suffice before the length of the list is $O(N^\epsilon)$ (the proof is identical to that of Lemma 4.2). Bounding the number of queries can be done similar to the proof of Lemma 4.3, which implies the bound on the total space. $\square$

Observing that the PRAM construction of Tarjan et al. [43] can be implemented in MPC gives the following lemma.

LEMMA 8.6. *Given a forest on n vertices, F , as a collection of undirected edges, the Euler tours for trees in F can be constructed in $O(1/\epsilon)$ rounds of MPC using $O(n)$ total space.*

Now, by applying forest connectivity, followed by the reduction from the tree rooting problem to list ranking described above we have the following theorem:

THEOREM 7. *Given a forest on n vertices, F , as a collection of undirected edges, the trees in F can be rooted in $O(1/\epsilon)$ rounds of AMPC w.h.p using $O(n \log n)$ total space w.h.p.*

Note that list-ranking allows us to efficiently implement algorithms based on the Euler tour of a tree. The idea is to first compute an Euler tour and apply list ranking from the root, which gives the distance from each edge to the root. After list ranking, the Euler tour can be represented as a sequence that we refer to as the *Euler sequence*, since for each edge its rank its position in an array representing the tour starting from the root. Computing properties, such as subtree sizes, or preorder numberings follow naturally from this representation of a tree.

Subtree Sizes. Given a rooted tree, T , we show how to compute the subtree sizes for each vertex in the tree. We first compute the Euler sequence of T , starting at the root. Next we assign a weight of 1 to each forward edge and a weight of 0 to each reverse edge (see the discussion on rooting trees above). Computing a prefix-sum over the Euler sequence with respect to the weights computes, for each edge, the number of forward edges preceding it. The subtree size at a given node, v , can be obtained by taking the difference of the values for the last reverse edge entering v and the first forward edge exiting v . Therefore, we have the following lemma:

LEMMA 8.7. *Computing the subtree sizes of a rooted tree can be performed in $O(1/\epsilon)$ rounds of AMPC w.h.p. using $O(n)$ total space w.h.p.*

Preorder Numbering. Computing a preorder numbering of a rooted tree can also be done using the Euler sequence and prefix-sums. The algorithm stores a weight of 1 for each forward edge and a weight of 0 for each reverse edge and computes the prefix-sum over the weighted Euler sequence with respect to the weights. The preorder number of a vertex is simply the number of forward edges preceding it. Therefore, we have the following lemma:

LEMMA 8.8. *Computing the preorder numbering of a rooted tree can be performed in $O(1/\epsilon)$ rounds of AMPC w.h.p. using $O(n)$ total space w.h.p.*

Subtree Minimum and Maximum. Note that the Euler tour-based technique described above to compute subtree sizes is not sufficient to compute the minimum and maximum over subtrees quickly (the technique depends on the function having an inverse, which precludes using it for min, max). The classic solution to this problem is the **range minimum query (RMQ)** data structure. The data structure is built over an input array of n elements. Queries take two indices, i, j and return the minimum value element in this range. Observe that after computing the Euler sequence, computing the minimum and maximum of a vertex v 's subtree can be reduced to a RMQ on the interval in the Euler tour representing v 's subtree. Using the fact that RMQ can be implemented efficiently in MPC (e.g., see Reference [2]) gives us the following lemma:

LEMMA 8.9. *A data structure to compute the minimum and maximum of a subtree can be computed in $O(1/\epsilon)$ rounds of AMPC. Each query to the data structure can be answered in $O(1)$ rounds. Preprocessing and storing the data structure uses $O(n)$ total space w.h.p.*

9 2-EDGE CONNECTIVITY

In this section, we show how our connectivity, spanning-forest and tree-based algorithms can be combined to obtain a 2-edge connectivity algorithm. A *biconnected component* of an undirected graph G is a maximal subgraph such that it remains connected even after removing any single vertex from it. A connected graph can be decomposed into a block-cut tree, which contains biconnected components, joined by vertices called *articulation points*. A graph is *k-edge connected* if it remains connected under the removal of any $k - 1$ edges. A *bridge* is an edge whose removal increases the number of connected components of G . In this section, we design a data structure that enables us

Sequentially, both problems can be solved using the Hopcroft–Tarjan algorithm [27]. The algorithm uses **depth-first search (DFS)** to identify the articulation points and bridges in $O(m + n)$ time. Although this sequential algorithm can be parallelized using parallel algorithms for DFS [1], the resulting algorithm requires many processors and is unlikely to lead to an efficient MPC or AMPC implementation. In the PRAM setting, Tarjan and Vishkin present the first work-efficient algorithm for biconnectivity and bridge-finding [43] (see also Maon et al. [33] and Ramachandran et al. [38] for biconnectivity algorithms based on computing open ear decompositions).

ALGORITHM 12: BC-LABELING($G=(V, E)$)

- (1) Let $(E_F, M) = \text{SPANNINGFOREST}(G)$. E_F is the set of spanning forest edges and M is the connectivity labeling of G .
- (2) Let $F = \text{ROOTFOREST}(E_F)$ and $PN = \text{PREORDERNUMBER}(F)$.
- (3) For each $v \in V$, compute $Low(v)$ and $High(v)$ and $Size(v)$ where $Low(v)$ and $High(v)$ are the minimum and maximum preorder numbers of all non-tree (u, w) edges where u is in v 's subtree, and where $Size(v)$ is the size of v 's subtree.
- (4) $critical = \{(u, p(u)) \in F \text{ s.t. } (u, p(u)) \text{ satisfies Equation (1)}\}$.
- (5) $L = \text{CONNECTIVITY}(G(V, E \setminus critical))$.
- (6) Output (L, F) .

Algorithm 12 describes our AMPC implementation of the Tarjan–Vishkin algorithm [43]. The algorithm first computes a spanning forest of G where the trees in the forest can be rooted

arbitrarily (Step 1). Next, it computes a rooted forest F by rooting each tree in the forest at an arbitrary vertex and computes PN , a preorder numbering of each tree (Step 2). It then computes for each $v \in V$ three quantities: $Low(v)$, $High(v)$, $Size(v)$, which are the minimum and maximum preorder numbers, respectively, of all non-tree (u, w) edges where u is a vertex in v 's subtree, and the size of each vertex's subtree. The initial values for $Low(v)$, $High(v)$, and $Size(v)$ are $PN(v)$, $PN(v)$, and 0 respectively.

The critical edges of G are tree edges $(v, p(v)) \in F$ where

$$PN(p(v)) \leq Low(v) \text{ and } High(v) \leq PN(p(v)) + Size(v). \quad (1)$$

The algorithm computes a labeling that can then be used to retrieve the bridges and articulation points in $O(1)$ queries as follows. We first delete the critical edges from the graph and compute the connectivity of this modified graph (Step 4). We follow Ben-David et al. and refer to this pair (L, F) as the *BC-labeling* of the graph [12].

Given the BC-labeling (L, F) , we can identify the bridges and articulation points as follows. The *head* of a component is the parent of the component in the spanning tree (the root is always a head). Each head defines a distinct biconnected component (there are at most n heads, since there can be at most n biconnected components). A non-root vertex v is an articulation point if it is the head of at least one component. The root of the forest is an articulation point if it is the head of two or more components, since these components were not connected by a non-tree edge in the graph with critical edges removed. A tree edge $(u, p(u))$ is a bridge if u 's component in L only contains u . Using this information we can compute 2-edge connected components of the graph by removing all bridges and running connectivity again.

Based on the description above, we have the following lemma about the correctness and cost of our BC-labeling algorithm:

LEMMA 9.1. *Given an undirected graph G , BC-labeling can be computed in $O(\log \log_{T/n} n)$ rounds of AMPC w.h.p. using $O(T)$ total space w.h.p. where the total space $T = \Omega(m + n)$.*

PROOF. The correctness of this algorithm follows from the proofs in Tarjan et al. [43] and Ben-David et al. [12].

The cost is proved by analyzing the steps of the algorithm. The spanning forest computation in Step 1 runs in $O(\log \log_{T/n} n)$ rounds of AMPC w.h.p. by Corollary 7.2. Step 2 runs in $O(1/\epsilon)$ rounds of AMPC w.h.p. and uses $O(n \log n)$ space w.h.p. by Theorem 7 and Lemma 8.8. Note that the space of the tree rooting algorithm can be reduced to $O(n)$ by applying the general graph connectivity algorithm (Algorithm 7) instead of the forest connectivity algorithm, at the expense of $O(\log \log_{T/n} n)$ rounds. Step 3, which computes Low , $High$, and $Size$ run in $O(1/\epsilon)$ rounds due to Lemmas 8.7 and 8.9. $\square$

Since we can use the BC-labeling to compute 2-edge connected components of the graph in the same round-complexity, we have the following theorem:

THEOREM 8. *2-edge connectivity can be computed in $O(\log \log_{T/n} n)$ rounds of AMPC w.h.p. where the total space $T = \Omega(m + n)$.*

10 OMITTED PROOFS

Here we prove the following lemma:

LEMMA 2.1. *Consider T balls, each having an integer weight between 0 and P . The total weight of all balls is T . Assume each of the balls is inserted uniformly at random to one of P bins. The bin for each ball is chosen independently of the choices for other balls and of the ball weight. Let $S = T/P$ and assume that $P = O(S^{1-\Omega(1)})$.*

Then, the total weight of balls in each bin is $O(S)$ with high probability.

PROOF. First we fix a certain bin b and show that the total weight of the balls in that bin is $O(S)$ with high probability. Then we use a union bound to show it for all bins simultaneously. For each ball $i \in \{1, \dots, T\}$ let X_i be a random variable that is 0 if ball i is not in bin b and be equal to the weight of ball i otherwise. Let $X = \sum_{i=1}^T X_i$ and note that $E[X] = \frac{T}{P} = S$. By the standard Bernstein's inequality we have

$$\begin{aligned} P(|X - E[X]| \leq E[X]) &\leq \exp\left(-\frac{E[X]^2/2}{\text{Var}(X) + PE[X]/3}\right) \\ &\leq \exp\left(-\frac{S^2}{2\text{Var}(X) + PS}\right), \end{aligned}$$

where $\text{Var}(X)$ is the variance of the random variable X . Next, we bound V . Let p_i be the weight of ball i . We have

$$\begin{aligned} \text{Var}(X) &= \sum_{i=1}^T \text{Var}(X_i) = \sum_{i=1}^T \left(\frac{1}{P} \left(1 - \frac{1}{P}\right) p_i^2\right) \\ &\geq \frac{1}{P} \sum_{i=1}^T p_i^2 \leq \frac{1}{P} \max(p_i) \sum_{i=1}^T p_i \\ &= \frac{1}{P} \cdot P \cdot T = T. \end{aligned}$$

Therefore we have

$$P(|X - E[X]| \leq E[X]) \leq \exp\left(-\frac{S^2}{2T + PS}\right) = \exp\left(-P^{\Omega(1)}\right),$$

which concludes the proof. $\square$

11 CONCLUSION AND OPEN PROBLEMS

In this work, we introduced and justified the Adaptive Massively Parallel Computation model, an extension of the Massively Parallel Computation model. Furthermore, we presented several new graph algorithms with much lower round complexities compared to the best-known algorithms in the MPC model, such as a constant round algorithm to find a maximal independent set and $O(\log \log_{m/n} n)$ round algorithms to find connectivity/minimum spanning tree among others problems. As future research, it is interesting to design algorithms for other problems such as vertex coloring, edge coloring and maximal matching in the AMPC model. Another open challenge is to develop (conditional or unconditional) hardness results in this model.

REFERENCES

- [1] Alok Aggarwal, Richard J. Anderson, and M.-Y. Kao. 1989. Parallel depth-first search in general directed graphs. In *Proceedings of the Annual ACM Symposium on Theory of Computing (STOC'89)*.
- [2] Alexandr Andoni, Zhao Song, Clifford Stein, Zhengyu Wang, and Peilin Zhong. 2018. Parallel graph connectivity in log diameter rounds. In *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science (FOCS'18)*, Mikkel Thorup (Ed.). IEEE Computer Society, 674–685. <https://doi.org/10.1109/FOCS.2018.00070>
- [3] Apache Software Foundation. [n.d.]. Beam. Retrieved from <https://beam.apache.org>.
- [4] Apache Software Foundation. [n.d.]. Giraph. Retrieved from <https://giraph.apache.org>.
- [5] Apache Software Foundation. [n.d.]. Hadoop. Retrieved from <https://hadoop.apache.org>.
- [6] InfiniBand Trade Association. [n.d.]. Retrieved from <https://www.infinibandta.org/>.
- [7] InfiniBand Trade Association. 2010. Supplement to InfiniBand Architecture Specification Volume 1 Release 1.2.2 Annex A16: RDMA over Converged Ethernet (RoCE).
- [8] MohammadHossein Bateni, Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Raimondas Kiveris, Silvio Lattanzi, and Vahab S. Mirrokni. 2017. Affinity clustering: Hierarchical clustering at scale. In *Advances in Neural Information Processing Systems 30: Proceedings of the Annual Conference on Neural Information Processing Systems*,

- Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). 6867–6877.
- [9] Paul Beame, Paraschos Koutris, and Dan Suciu. 2017. Communication steps for parallel query processing. *J. ACM* 64, 6 (2017), 40:1–40:58. <https://doi.org/10.1145/3125644>
 - [10] Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Łacki, and Vahab Mirrokni. 2019. Near-optimal massively parallel graph connectivity. In *Proceedings of the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS'19)*. 1615–1636.
 - [11] Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Łacki, Vahab S. Mirrokni, and Warren Schudy. 2020. Parallel graph algorithms in constant adaptive rounds: Theory meets practice. *Proc. VLDB Endow.* 13, 13 (2020), 3588–3602.
 - [12] Naama Ben-David, Guy E. Blelloch, Jeremy T. Fineman, Phillip B. Gibbons, Yan Gu, Charles McGuffey, and Julian Shun. 2018. Implicit decomposition for write-efficient connectivity algorithms. In *Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS'18)*.
 - [13] Mark S. Birrittella, Mark Debbage, Ram Huggahalli, James Kunz, Tom Lovett, Todd Rimmer, Keith D. Underwood, and Robert C. Zak. 2015. Intel®Omni-path architecture: Enabling scalable, high performance fabrics. In *Proceedings of the 2015 IEEE 23rd Annual Symposium on High-Performance Interconnects (HOTI'15)*. IEEE Computer Society, Los Alamitos, CA, 1–9. <https://doi.org/10.1109/HOTI.2015.22>
 - [14] Guy E. Blelloch, Jeremy T. Fineman, and Julian Shun. 2012. Greedy sequential maximal independent set and matching are parallel on average. In *Proceedings of the 24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'12)*. 308–317. <https://doi.org/10.1145/2312005.2312058>
 - [15] Guy E. Blelloch and Bruce M. Maggs. 2010. Parallel Algorithms. In *Algorithms and Theory of Computation Handbook*. Chapman & Hall/CRC, 25–25. <http://dl.acm.org/citation.cfm?id=1882723.1882748>.
 - [16] Craig Chambers, Ashish Raniwala, Frances Perry, Stephen Adams, Robert Henry, Robert Bradshaw, and Nathan. 2010. FlumeJava: Easy, efficient data-parallel pipelines. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'10)*. ACM, New York, NY, 363–375. <http://dl.acm.org/citation.cfm?id=1806638>.
 - [17] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Michael Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. 2008. Bigtable: A distributed storage system for structured data. *ACM Trans. Comput. Syst.* 26, 2 (2008), 4:1–4:26. <https://doi.org/10.1145/1365815.1365816>
 - [18] Moses Charikar, Weiyun Ma, and Li-Yang Tan. 2020. Unconditional lower bounds for adaptive massively parallel computation. In *Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures, Virtual Event (SPAA'20)*. 141–151. <https://doi.org/10.1145/3350755.3400230>
 - [19] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms (3rd edition)*. MIT Press.
 - [20] Intel Corporation. [n.d.]. Intel Omni-Path Architecture. Retrieved from <https://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-driving-exascale-computing.html>.
 - [21] Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: Simplified data processing on large clusters. *Commun. ACM* 51, 1 (2008), 107–113. <https://doi.org/10.1145/1327452.1327492>
 - [22] Aleksandar Dragojevic, Dushyanth Narayanan, and Miguel Castro. 2017. RDMA reads: To use or not to use? *IEEE Data Eng. Bull.* 40, 1 (2017), 3–14.
 - [23] Aleksandar Dragojević, Dushyanth Narayanan, Miguel Castro, and Orion Hodson. 2014. FaRM: Fast remote memory. In *Proceedings of the 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI'14)*. 401–414.
 - [24] Mohsen Ghaffari, Fabian Kuhn, and Jara Uitto. 2019. Conditional hardness results for massively parallel computation from distributed lower bounds. In *Proceedings of the 60th IEEE Annual Symposium on Foundations of Computer Science (FOCS'19)*. 1650–1663. <https://doi.org/10.1109/FOCS.2019.00097>
 - [25] Mohsen Ghaffari and Jara Uitto. 2019. Sparsifying distributed algorithms with ramifications in massively parallel computation and centralized local computation. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'19)*, Timothy M. Chan (Ed.). SIAM, 1636–1653. <https://doi.org/10.1137/1.9781611975482.99>
 - [26] Michael T. Goodrich, Nodari Sitchinava, and Qin Zhang. 2011. Sorting, searching, and simulation in the mapreduce framework. In *Proceedings of the 22nd International Symposium on Algorithms and Computation (ISAAC'11)*, Takao Asano, Shin-Ichi Nakano, Yoshio Okamoto, and Osamu Watanabe (Eds.), Lecture Notes in Computer Science, Vol. 7074. Springer, 374–383. https://doi.org/10.1007/978-3-642-25591-5_39
 - [27] John Hopcroft and Robert Tarjan. 1973. Algorithm 447: Efficient algorithms for graph manipulation. *Commun. ACM* (1973).
 - [28] Anuj Kalia, Michael Kaminsky, and David G. Andersen. 2014. Using RDMA efficiently for key-value services. In *Proceedings of the 2014 ACM Conference on SIGCOMM (SIGCOMM'14)*. ACM, New York, NY, 295–306. <https://doi.org/10.1145/2619239.2626299>

- [29] 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'10)*, Moses Charikar (Ed.). SIAM, 938–948. <https://doi.org/10.1137/1.9781611973075.76>
- [30] Raimondas Kiveris, Silvio Lattanzi, Vahab S. Mirrokni, Vibhor Rastogi, and Sergei Vassilvitskii. 2014. Connected components in mapreduce and beyond. In *Proceedings of the ACM Symposium on Cloud Computing*, Ed Lazowska, Doug Terry, Remzi H. Arpaci-Dusseau, and Johannes Gehrke (Eds.). ACM, 18:1–18:13. <https://doi.org/10.1145/2670979.2670997>
- [31] Hyeontaek Lim, Dongsu Han, David G. Andersen, and Michael Kaminsky. 2014. MICA: A holistic approach to fast in-memory key-value storage. In *Proceedings of the 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI'14)*. USENIX Association, 429–444. <https://www.usenix.org/conference/nsdi14/technical-sessions/presentation/lim>.
- [32] 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 ACM SIGMOD International Conference on Management of Data (SIGMOD'10)*, Ahmed K. Elmagarmid and Divyakant Agrawal (Eds.). ACM, 135–146. <https://doi.org/10.1145/1807167.1807184>
- [33] Yael Maon, Baruch Schieber, and Uzi Vishkin. 1986. Parallel ear decomposition search (EDS) and st-numbering in graphs. *Theor. Comput. Sci.* 47 (1986), 277–298.
- [34] Christopher Mitchell, Yifeng Geng, and Jinyang Li. 2013. Using one-sided RDMA reads to build a fast, CPU-efficient Key-value Store. In *Proceedings of the 2013 USENIX Conference on Annual Technical Conference (USENIX ATC'13)*. USENIX Association, Berkeley, CA, 103–114.
- [35] Rajeev Motwani and Prabhakar Raghavan. 1995. *Randomized Algorithms*. Cambridge University Press.
- [36] Huy N. Nguyen and Krzysztof Onak. 2008. Constant-time approximation algorithms via local improvements. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS'08)*. 327–336. <https://doi.org/10.1109/FOCS.2008.81>
- [37] John Ousterhout, Arjun Gopalan, Ashish Gupta, Ankita Kejriwal, Collin Lee, Behnam Montazeri, Diego Ongaro, Seo Jin Park, Henry Qin, Mendel Rosenblum, et al. 2015. The RAMCloud storage system. *ACM Trans. Comput. Syst.* 33, 3 (2015), 7.
- [38] Vijaya Ramachandran. 1993. Parallel open ear decomposition with applications to graph biconnectivity and triconnectivity. In *Synthesis of Parallel Algorithms*.
- [39] Abhiram Gorakhanath Ranade. 1989. *Fluent Parallel Computation*. Ph.D. Dissertation. New Haven, CT.
- [40] Tim Roughgarden, Sergei Vassilvitskii, and Joshua R. Wang. 2016. Shuffles and Circuits (on lower bounds for modern parallel computation). In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'16)*, Christian Scheideler and Seth Gilbert (Eds.). ACM, 1–12. <https://doi.org/10.1145/2935764.2935799>
- [41] Peter Sanders. [n.d.]. On the competitive analysis of randomized static load balancing. In *Proceedings of the 1st Workshop on Randomized Parallel Algorithms*. 1–21.
- [42] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, and Robert Chansler. 2010. The hadoop distributed file system. In *Proceedings of the IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST'12)*, Mohammed G. Khatib, Xubin He, and Michael Factor (Eds.). IEEE Computer Society, 1–10. <https://doi.org/10.1109/MSST.2010.5496972>
- [43] Robert E. Tarjan and Uzi Vishkin. 1985. An efficient parallel biconnectivity algorithm. *SIAM J. Comput.* (1985).
- [44] Leslie G. Valiant. 1990. A bridging model for parallel computation. *Commun. ACM* 33, 8 (1990), 103–111.
- [45] Aad W. van der Vaart and Jon A. Wellner. [n.d.]. *Weak convergence and empirical processes with applications to statistics*. Springer Science & Business Media.
- [46] Grigory Yaroslavltssev and Adithya Vadapalli. 2018. Massively parallel algorithms and hardness for single-linkage clustering under ℓ_p -distances. In *Proceedings of the 35th International Conference on Machine Learning (ICML'18)*, Jennifer G. Dy and Andreas Krause (Eds.), JMLR Workshop and Conference Proceedings, Vol. 80. 5596–5605.
- [47] Yuichi Yoshida, Masaki Yamamoto, and Hiro Ito. 2009. An improved constant-time approximation algorithm for maximum matchings. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC'09)*, Michael Mitzenmacher (Ed.). ACM, 225–234. <https://doi.org/10.1145/1536414.1536447>
- [48] Matei Zaharia, Reynold S. Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman, Michael J. Franklin, Ali Ghodsi, Joseph Gonzalez, Scott Shenker, and Ion Stoica. 2016. Apache Spark: A unified engine for big data processing. *Commun. ACM* 59, 11 (2016), 56–65. <https://doi.org/10.1145/2934664>
- [49] Erfan Zamanian, Carsten Binnig, Tim Harris, and Tim Kraska. 2017. The end of a myth: Distributed transactions can scale. *Proc. VLDB Endow.* 10, 6 (2017), 685–696.

Received December 2019; revised September 2020; accepted October 2020