

Massively Parallel Computation via Remote Memory Access

Soheil Behnezhad
University of Maryland
College Park, USA
soheil@cs.umd.edu

Laxman Dhulipala
CMU
Pittsburgh, USA
ldhulipa@andrew.cmu.edu

Hossein Esfandiari
Google
New York, USA
esfandiari@google.com

Jakub Łącki
Google
New York, USA
jłacki@google.com

Vahab Mirrokni
Google
New York, USA
mirrokni@google.com

Warren Schudy
Google
New York, USA
wschudy@google.com

ABSTRACT

We introduce the *Adaptive Massively Parallel Computation* (AMPC) model, which is an extension of the widely popular *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 [9, 28] 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, which is an *exponential* improvement upon the best known algorithms in the MPC model with sublinear space per machine. Our results imply that the 2-CYCLE conjecture, the most popular hardness conjecture in the MPC model, does not hold in the AMPC model.

CCS CONCEPTS

• **Theory of computation** → **Distributed computing models; Distributed algorithms; Graph algorithms analysis;**

KEYWORDS

distributed algorithms, graph algorithms, massively parallel computation, remote memory access

ACM Reference Format:

Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Łącki, Vahab Mirrokni, and Warren Schudy. 2019. Massively Parallel Computation via Remote Memory Access. In *31st ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '19), June 22–24, 2019, Phoenix, AZ, USA*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3323165.3323208>



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

SPAA '19, June 22–24, 2019, Phoenix, AZ, USA

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6184-2/19/06.

<https://doi.org/10.1145/3323165.3323208>

1 INTRODUCTION

The MapReduce computation framework was introduced in 2004 by Dean and Ghemawat [20]. 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 [45], Hadoop [5], FlumeJava [16], Beam [3], Pregel [31], and Giraph [4].

These frameworks share many similarities, especially from a theoretical point of view. For this reason, the *Massively Parallel Computation* (MPC) model [10, 24, 27] 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 RDMA and RPCs 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 [6, 29, 38, 43] 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 remote direct-memory access (RDMA) without requiring a synchronized round of communication.

In this work, 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 paper we introduce the *Adaptive Massively Parallel Computation* model (AMPC), 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.

Problem	AMPC	MPC
Connectivity	$O(\log \log_{m/n} n)$	$O(\log D \cdot \log \log_{m/n} n)$ [2]
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	$O(1)$	$\tilde{O}(\sqrt{\log n})$ [23]
2-CYCLE	$O(1)$	$O(\log n)$
Forest Connectivity	$O(1)$	$O(\log D \cdot \log \log_{m/n} n)$ [2]

Figure 1: Round complexities of our graph algorithms in the AMPC model, compared to the best known results in the MPC model. D denotes the diameter of the input graph. We consider the setting where the space per machine is sublinear in the number of vertices of the graph, that is $S = O(n^\epsilon)$.

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. Hence, under the conjecture, the AMPC model is strictly stronger than the MPC model.

Our model is inspired by two previous papers that combine the use of distributed hash table service (DHT) and MapReduce [9, 28]. The first paper [28] 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 [9] 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 paper 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). Due to space constraints we defer some of the proofs and some implementation details to the full version.

2 MODEL

In this section, we formalize the *Adaptive Massively Parallel Computation* model (AMPC).

In the AMPC model, we have an input of size N^1 , which is processed by P machines, each with space S . 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 data stores (DDS) that we denote by $\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2, \dots$. For the description, it is convenient to assume that all DDS-es provide key-value store semantics, that they store a collection of key-value pairs: given a key a DDS returns the corresponding value. We require each key-value pair to have constant size (i.e. both key and value consist of a constant number of words). The input data is stored in \mathcal{D}_0 and uses a set of keys known to all machines (e.g. consecutive integers).

If $k > 1$ key-value pairs stored in the DDS have the same key x , the individual values can be accessed by querying for keys $(x, 1), \dots, (x, k)$. Note that the indices from 1 to k are assigned arbitrarily. Querying for a key that does not occur in the DDS, results in an empty response.

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 up to $O(S)$ reads (henceforth called *queries*) and $O(S)$ writes, and perform arbitrary computation². Each query and write refers to querying for/writing a single key-value pair. The amount of *communication* that a machine performs per round is equal to the total number of queries and writes.

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 $x \in X$, \mathcal{D}_{i-1} stores a key-value pair $(x, g(x))$, then in round i a machine can compute $g^k(y)$ in a single round, provided that $k = O(S)$.

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 DDS. In the following round, each machine reads all key-value pairs keyed by its id. The limits on the number of queries and writes

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 paper 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.

in the AMPC model are direct counterparts of the communication limits of the MPC model.

Due to known simulations of PRAM algorithms by MPC [24, 27], 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 [37]. The simulations can be done using $O(1)$ rounds per PRAM step, and total space proportional to the number of processors.

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 remote direct memory-access (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 [7], Omnipath [13, 19] and RoCE [8] being widely available. RDMA has been used in recent work on key-value storage systems [22, 26, 30, 33], database systems [46], and other related systems and studies [21, 35]. The latency for a remote memory access in recent RDMA systems is around 1–3 microseconds under load, making it about 20x more costly than a main-memory access [21].

Second, we observe that the per-machine latency can be further reduced by using the *parallel slackness* technique [42]. 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 that we present in this paper.

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 [20, 31, 40, 45]. 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 DDS 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 DDS in each round, but the machines are otherwise not restricted in accessing the DDS. In particular, all machines may be accessing the same key-value pair, which could potentially lead to contention. We observe that under a few natural assumptions, this should not be an issue. Let us first discuss the assumptions.

- (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 [20, 31]. To make sure that the amount of this communication is bounded by $O(S)$, we must have $P = O(S)$.
- (2) The DDS 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 DDS. 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 DDS needs to answer $O(S)$ queries with high probability. 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.

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 DDS may have changed since the round has started. Moreover, if writes were allowed, 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 paper 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 $T = n^{1+\Omega(1)}$.

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 which rely on AMPC model features. The remaining parts can be implemented in MPC model using standard primitives, such as sorting, duplicate removal, etc.

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. In order to achieve this, it suffices to store in the DDS 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) Let $N = |V|$
- (2) 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.
- (3) Return the graph $(M, \{xl_x \mid x \in M\} \cup \{xr_x \mid x \in M\})$.

Algorithm 2. $2\text{-CYCLE}(G = (V, E))$

- (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.

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 which can be obtained from G by contracting edges. The length of each cycle in G' is $O(n^\epsilon)$.*

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 .*

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 $O(T)$ 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 $O(1/\epsilon)$ 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 by [34, 44] 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 $O(1/\epsilon)$ 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. [44] 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. $\text{TRUNCATEDQUERY}(v, \pi, c)$

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 $\text{is-in-MIS}(v) = \text{true}$ 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 $\text{is-in-MIS}(v) = \text{false}$ 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 $O(1/\epsilon)$ 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 $\text{is-in-MIS}(v) \neq \text{unknown}$.

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 $\text{TRUNCATEDQUERY}(v, \pi, c)$. If v is not settled by the end of iteration r , 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.*

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 of [14] (see the paragraph below) by $D = O(\log^2 n)$. Thus, a total of $\Omega\left(\frac{n^\epsilon}{\log^2 n} \cdot n^{(r-1)\epsilon/2}\right) \gg n^{r\epsilon/2}$ queries should not be sufficient to settle the status of v . This is a contradiction since we assumed $q_v \leq n^{r\epsilon/2}$.

Bounding the recursion depth. The analysis of [14] bounds the round complexity of a parallel implementation of the greedy (lexicographically first) MIS algorithm by $O(\log^2 n)$. While applying this result as a black-box does not bound the depth of recursions, here we argue that with minor modifications, it also bounds the recursion depth by $O(\log^2 n)$. The analysis of [14] divides the vertex set into disjoint layers V_1, V_2, \dots, V_k where $k = O(\log n)$ such that the vertices in V_i have higher priority of joining MIS than all those in V_j for $j > i$. The key property is that after finding the vertices of the lexicographically first MIS among V_1, \dots, V_i for any i , and removing their neighbors, the remaining vertices in V_{i+1} have logarithmic degree within their partition and thus, any path going through vertices of increasing priorities has length $O(\log n)$. Collecting all these paths of all layers, this gives a $O(\log^2 n)$ round complexity for the parallel implementation of the algorithm.

In order to adapt this to our setting, suppose that the recursion depth of a vertex is larger than $\beta \log^2 n$ for some large enough constant β . This means that in one of the layers defined by [14], say V_i , the query process has to visit more than $\alpha \log n$ vertices for some sufficiently large constant α . Note that by the analysis of [14], there should be a vertex in layers V_1, \dots, V_{i-1} that destroys this path by arguments of [14]. Let v be a vertex that has to be destroyed. This means that v has to have a neighbor in lower layers which joins the MIS. The crucial observation, here, is that since the query process is greedy, v will immediately go to lower levels and will be eventually deleted from the graph before visiting any other vertex in layer V_i . Thus, the arguments of [14] still hold and the recursion depth is bounded by $O(\log^2 n)$.

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 $T = n^{1+\Omega(1)}$, 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, 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 DDS. 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 $n^{1-\epsilon}d^2$, the maximum number of queries a machine makes is $O(n^\epsilon)$ w.h.p.*

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 $\log D$.

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. In order 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 of [11] 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 ([11]). *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 $T = \Omega(m + n)$.*

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 [18]. 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. We first decrease the number of vertices in the graph if $m = o(n \log^2 n)$. Next, we 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 $O(1/\epsilon)$ rounds.

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

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. using $O(T)$ total space.*

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 on forests in the AMPC model. 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 of Tarjan *et al.* [41] can be implemented in AMPC implies that this

reduction to the cycle connectivity problem can be done in $O(1)$ rounds of AMPC.

Next, we address the cycle connectivity problem.

Cycle connectivity. The *cycle connectivity problem* is to compute the connected components of a graph $G = (V, E)$ which contains a set of disjoint cycles. Our cycle connectivity algorithm, Algorithm 10, works as follows. We call $O(1/\epsilon)$ iterations of SHRINK (the algorithm we used in the 2-CYCLE problem) which reduces the size of the largest cycle to $O(n^\epsilon)$ with high probability. Then, we switch to a different algorithm which 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.*

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.*

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 in order 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 [39], it does not provide 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.*

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 using $T = O(n \log n)$ total space.*

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 [25]. 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 [41] (see also Maon *et al.* [32] and Ramachandran *et al.* [36] for biconnectivity algorithms based on computing open ear decompositions).

Algorithm 11. 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 11 describes our AMPC implementation of the Tarjan-Vishkin algorithm [41]. 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 quickly 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.*

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 6. *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 CONCLUSION AND OPEN PROBLEMS

We have introduced and justified the Adaptive Massively Parallel Computation model, an extension of the widely popular 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 maximum 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 *STOC*.
- [2] Alexandr Andoni, Zhao Song, Clifford Stein, Zhengyu Wang, and Peilin Zhong. 2018. Parallel Graph Connectivity in Log Diameter Rounds. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, Mikkel Thorup (Ed.). IEEE Computer Society, 674–685. <https://doi.org/10.1109/FOCS.2018.00070>
- [3] Apache Software Foundation. Beam. <https://beam.apache.org/>.
- [4] Apache Software Foundation. Giraph. <https://giraph.apache.org/>.
- [5] Apache Software Foundation. Hadoop. <https://hadoop.apache.org/>.
- [6] Sepehr Assadi, Xiaorui Sun, and Omri Weinstein. 2018. Massively Parallel Algorithms for Finding Well-Connected Components in Sparse Graphs. *CoRR* abs/1805.02974 (2018). arXiv:1805.02974 <http://arxiv.org/abs/1805.02974>
- [7] InfiniBand Trade Association. <https://www.infinibandta.org/>

- [8] InfiniBand Trade Association. 2010. Supplement to InfiniBand Architecture Specification Volume 1 Release 1.2.2 Annex A16: RDMA over Converged Ethernet (RoCE). (2010).
- [9] MohammadHossein Bateni, Soheil Behnezhad, Mahsa Derakhshan, Mohammad-Taghi Hajiaghayi, Raimondas Kiveris, Silvio Lattanzi, and Vahab S. Mirrokni. 2017. Affinity Clustering: Hierarchical Clustering at Scale. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.), 6867–6877. <http://papers.nips.cc/paper/7262-affinity-clustering-hierarchical-clustering-at-scale>
- [10] 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>
- [11] Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakob Lacki, and Vahab Mirrokni. 2019. Near-Optimal Massively Parallel Graph Connectivity. (2019). manuscript.
- [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 *IPDPS*.
- [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, Washington, DC, USA, 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 *24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '12, Pittsburgh, PA, USA, June 25-27, 2012*. 308–317. <https://doi.org/10.1145/2312005.2312058>
- [15] Guy E. Blelloch and Bruce M. Maggs. 2010. Algorithms and Theory of Computation Handbook. Chapman & Hall/CRC, Chapter Parallel Algorithms, 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 *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. 2 Penn Plaza, Suite 701 New York, NY 10121-0701, 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] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms (3rd edition)*. MIT Press.
- [19] Intel Corporation. Intel Omni-Path Architecture. <https://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-driving-exascale-computing.html>
- [20] 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>
- [21] Aleksandar Dragojević, Dushyanth Narayanan, and Miguel Castro. 2017. RDMA Reads: To Use or Not to Use? *IEEE Data Eng. Bull.* 40, 1 (2017), 3–14.
- [22] Aleksandar Dragojević, Dushyanth Narayanan, Miguel Castro, and Orion Hodson. 2014. FaRM: Fast remote memory. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. 401–414.
- [23] Mohsen Ghaffari and Jara Uitto. 2019. Sparsifying Distributed Algorithms with Ramifications in Massively Parallel Computation and Centralized Local Computation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, Timothy M. Chan (Ed.). SIAM, 1636–1653. <https://doi.org/10.1137/1.9781611975482.99>
- [24] Michael T. Goodrich, Nodari Sitchinava, and Qin Zhang. 2011. Sorting, Searching, and Simulation in the MapReduce Framework. In *Algorithms and Computation - 22nd International Symposium, ISAAC 2011, Yokohama, Japan, December 5-8, 2011. Proceedings (Lecture Notes in Computer Science)*, Takao Asano, Shin-Ichi Nakano, Yoshio Okamoto, and Osamu Watanabe (Eds.), Vol. 7074. Springer, 374–383. https://doi.org/10.1007/978-3-642-25591-5_39
- [25] John Hopcroft and Robert Tarjan. 1973. Algorithm 447: efficient algorithms for graph manipulation. *Commun. ACM* (1973).
- [26] 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, USA, 295–306. <https://doi.org/10.1145/2619239.2626299>
- [27] Howard J. Karloff, Siddharth Suri, and Sergei Vassilvitskii. 2010. A Model of Computation for MapReduce. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, Moses Charikar (Ed.). SIAM, 938–948. <https://doi.org/10.1137/1.9781611973075.76>
- [28] 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, Seattle, WA, USA, November 3-5, 2014*, Ed Lazowska, Doug Terry, Remzi H. Arpaci-Dusseau, and Johannes Gehrke (Eds.). ACM, 18:1–18:13. <https://doi.org/10.1145/2670979.2670997>
- [29] Jakob Lacki, Vahab S. Mirrokni, and Michal Włodarczyk. 2018. Connected Components at Scale via Local Contractions. *CoRR* abs/1807.10727 (2018). arXiv:1807.10727 <http://arxiv.org/abs/1807.10727>
- [30] Hyeontaek Lim, Dongsu Han, David G. Andersen, and Michael Kaminsky. 2014. MICA: A Holistic Approach to Fast In-Memory Key-Value Storage. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. USENIX Association, Seattle, WA, 429–444. <https://www.usenix.org/conference/nsdi14/technical-sessions/presentation/lim>
- [31] 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 2010, Indianapolis, Indiana, USA, June 6-10, 2010*, Ahmed K. Elmagarmid and Divyakant Agrawal (Eds.). ACM, 135–146. <https://doi.org/10.1145/1807167.1807184>
- [32] Yael Maon, Baruch Schieber, and Uzi Vishkin. 1986. Parallel ear decomposition search (EDS) and st-numbering in graphs. *Theoretical Computer Science* 47 (1986).
- [33] 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, USA, 103–114. <http://dl.acm.org/citation.cfm?id=2535461.2535475>
- [34] 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)*. 327–336. <https://doi.org/10.1109/FOCS.2008.81>
- [35] 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 Transactions on Computer Systems (TOCS)* 33, 3 (2015), 7.
- [36] Vijaya Ramachandran. 1993. Parallel open ear decomposition with applications to graph biconnectivity and triconnectivity. In *Synthesis of Parallel Algorithms*.
- [37] Abhiram Gorakhanath Ranade. 1989. *Fluent Parallel Computation*. Ph.D. Dissertation. New Haven, CT, USA. AAI9010675.
- [38] 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 2016, Asilomar State Beach/Pacific Grove, CA, USA, July 11-13, 2016*, Christian Scheideler and Seth Gilbert (Eds.). ACM, 1–12. <https://doi.org/10.1145/2935764.2935799>
- [39] Peter Sanders. 1996. On the competitive analysis of randomized static load balancing. In *1st Workshop on Randomized Parallel Algorithms*.
- [40] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, and Robert Chansler. 2010. The Hadoop Distributed File System. In *IEEE 26th Symposium on Mass Storage Systems and Technologies, MSST 2012, Lake Tahoe, Nevada, USA, May 3-7, 2010*, Mohammed G. Khatib, Xubin He, and Michael Factor (Eds.). IEEE Computer Society, 1–10. <https://doi.org/10.1109/MSST.2010.5496972>
- [41] Robert E Tarjan and Uzi Vishkin. 1985. An efficient parallel biconnectivity algorithm. *SIAM J. Comput.* (1985).
- [42] Leslie G Valiant. 1990. A bridging model for parallel computation. *Commun. ACM* 33, 8 (1990), 103–111.
- [43] Grigory Yaroslavtsev 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 2018, Stockholm, Sweden, July 10-15, 2018 (JMLR Workshop and Conference Proceedings)*, Jennifer G. Dy and Andreas Krause (Eds.), Vol. 80. JMLR.org, 5596–5605. <http://proceedings.mlr.press/v80/yaroslavtsev18a.html>
- [44] 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 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, Michael Mitzenmacher (Ed.). ACM, 225–234. <https://doi.org/10.1145/1536414.1536447>
- [45] 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>
- [46] Erfan Zamanian, Carsten Binnig, Tim Harris, and Tim Kraska. 2017. The end of a myth: Distributed transactions can scale. *Proceedings of the VLDB Endowment* 10, 6 (2017), 685–696.