

Connected Components on a PRAM in Log Diameter Time

S. Cliff Liu
Princeton University
sixuel@cs.princeton.edu

Robert E. Tarjan
Princeton University
ret@cs.princeton.edu

Peilin Zhong
Columbia University
peilin.zhong@columbia.edu

March 3, 2020

Abstract

We present an $O(\log d + \log \log_{m/n} n)$ -time randomized PRAM algorithm for computing the connected components of an n -vertex, m -edge undirected graph with maximum component diameter d . The algorithm runs on an ARBITRARY CRCW (concurrent-read, concurrent-write with arbitrary write resolution) PRAM using $O(m)$ processors. The time bound holds with good probability.¹

Our algorithm is based on the breakthrough results of Andoni et al. [FOCS'18] and Behnezhad et al. [FOCS'19]. Their algorithms run on the more powerful MPC model and rely on sorting and computing prefix sums in $O(1)$ time, tasks that take $\Omega(\log n / \log \log n)$ time on a CRCW PRAM with $\text{poly}(n)$ processors. Our simpler algorithm uses limited-collision hashing and does not sort or do prefix sums. It matches the time and space bounds of the algorithm of Behnezhad et al., who improved the time bound of Andoni et al.

It is widely believed that the larger private memory per processor and unbounded local computation of the MPC model admit algorithms faster than that on a PRAM. Our result suggests that such additional power might not be necessary, at least for fundamental graph problems like connected components and spanning forest.

1 Introduction

Computing the connected components of an undirected graph is a fundamental problem in algorithmic graph theory, with many applications. Using graph search [Tar72], one can find the connected components of an n -vertex, m -edge graph in $O(m)$ time, which is best possible for a sequential algorithm. But linear time is not fast enough for big-data applications in which the problem graph is of internet scale or even bigger. To find the components of such large graphs in practice requires the use of concurrency.

Beginning in the 1970's, theoreticians developed a series of more-and-more efficient concurrent algorithms. Their model of computation was some variant of the PRAM (parallel random access machine) model. Shiloach and Vishkin [SV82] gave an $O(\log n)$ -time PRAM algorithm in 1982. The algorithm was later simplified by Awerbuch and Shiloach [AS87]. These algorithms are deterministic and run on an ARBITRARY CRCW PRAM with $O(m)$ processors. There are simpler algorithms and algorithms that do less work but use randomization. Gazit [Gaz91] combined an elegant randomized algorithm of Reif [Rei84] with graph size reduction to obtain an $O(\log n)$ -time, $O(m / \log n)$ -processor CRCW PRAM algorithm. This line of work culminated in the $O(\log n)$ -time, $O(m / \log n)$ -processor EREW (exclusive-read, exclusive-write) PRAM algorithms of Halperin and Zwick [HZ96, HZ01], the second of which computes spanning trees of the components as well as the components themselves. On an EREW PRAM, finding connected components takes $\Omega(\log n)$ time [CDR86], so these algorithms minimize both time and work. The $\Omega(\log n)$ lower bound

¹To simplify the statements in this paper we assume $m/2 \geq n \geq 2$ and $d \geq 1$. *With good probability* means with probability at least $1 - 1/\text{poly}((m \log n)/n)$; *with high probability* means with probability at least $1 - 1/\text{poly}(n)$.

also holds for the CREW (concurrent-read, exclusive-write) PRAM with randomization, a model slightly weaker than the CRCW PRAM [DKR94]. For the PRIORITY CRCW PRAM (write resolution by processor priority), a time bound of $\Omega(\log n / \log \log n)$ holds if there are $\text{poly}(n)$ processors and unbounded space, or if there is $\text{poly}(n)$ space and any number of processors [BH89].

The Halperin-Zwick algorithms use sophisticated techniques. Practitioners charged with actually finding the connected components of huge graphs have implemented much simpler algorithms. Indeed, such simple algorithms often perform well in practice [GKP94, Gre94, HRD97, SDB14, SRT18]. The computational power of current platforms and the characteristics of the problem graphs may partially explain such good performance. Current parallel computing platforms such as MapReduce, Hadoop, Spark, and others have capabilities significantly beyond those modeled by the PRAM [DG08]. A more powerful model, the MPC (massively parallel computing) model [BKS17] is intended to capture these capabilities. In this model, each processor can have $\text{poly}(n)$ (typically sublinear in n) private memory, and the local computational power is unbounded. A PRAM algorithm can usually be simulated on an MPC with asymptotically the same round complexity, and it is widely believed that the MPC model admits algorithms faster than the PRAM model [KSV10, GSZ11, ASW19]. On the MPC model, and indeed on the weaker COMBINING CRCW PRAM model, there are very simple, practical algorithms that run in $O(\log n)$ time [LT19]. Furthermore, many graphs in applications have components of small diameter, perhaps poly-logarithmic in n .

These observations lead to the question of whether one can find connected components faster on graphs of small diameter, perhaps by exploiting the power of the MPC model. Andoni et al. [ASS⁺18] answered this question “yes” by giving an MPC algorithm that finds connected components in $O(\log d \log \log_{m/n} n)$ time, where d is the largest diameter of a component. Very recently, this time bound was improved to $O(\log d + \log \log_{m/n} n)$ by Behnezhad et al. [BDE⁺19]. Both of these algorithms are complicated and use the extra power of the MPC model, in particular, the ability to sort and compute prefix sums in $O(1)$ communication rounds. These operations require $\Omega(\log n / \log \log n)$ time on a CRCW PRAM with $\text{poly}(n)$ processors [BH89].² These results left open the following fundamental problem in theory:

Is it possible to break the $\log n$ time barrier for connected components of small-diameter graphs on a PRAM (without the additional power of an MPC)?

In this paper we give a positive answer by presenting an ARBITRARY CRCW PRAM algorithm that runs in $O(\log d + \log \log_{m/n} n)$ time, matching the round complexity of the MPC algorithm by Behnezhad et al. [BDE⁺19]. Our algorithm uses $O(m)$ processors and space, and thus is space-optimal and nearly work-efficient. In contrast to the MPC algorithm, which uses several powerful primitives, we use only hashing and other simple data structures. Our hashing-based approach also applies to the work of Andoni et al. [ASS⁺18], giving much simpler algorithms for connected components and spanning forest, which should be preferable in practice.

1.1 Computation Models and Main Results

Our main model of computation is the ARBITRARY CRCW PRAM [Vis83]. It consists of a set of processors, each of which has a constant number of cells (words) as the private memory, and a large common memory. The processors run synchronously. In one step, a processor can read a cell in common memory, write to a cell in common memory, or do a constant amount of local computation. Any number of processors can read from or write to the same common memory cell concurrently. If more than one processor writes to the same memory cell at the same time, an arbitrary one succeeds.

Our main results are stated below:

²Behnezhad et al. also consider the *multiprefix* CRCW PRAM, in which prefix sum (and other primitives) can be computed in $O(1)$ time and $O(m)$ work. A direct simulation of this model on a PRIORITY CRCW PRAM or weaker model would suffer an $\Omega(\log n / \log \log n)$ factor in both time and work, compared to our result.

Theorem 1 (Connected Components). *There is an ARBITRARY CRCW PRAM algorithm using $O(m)$ processors that computes the connected components of any given graph. With probability $1 - 1/\text{poly}((m \log n)/n)$, it runs in $O(\log d \log \log_{m/n} n)$ time.*

The algorithm of Theorem 1 can be extended to computing a spanning forest (a set of spanning trees of the components) with the same asymptotic efficiency:

Theorem 2 (Spanning Forest). *There is an ARBITRARY CRCW PRAM algorithm using $O(m)$ processors that computes the spanning forest of any given graph. With probability $1 - 1/\text{poly}((m \log n)/n)$, it runs in $O(\log d \log \log_{m/n} n)$ time.*

Using the above algorithms as bases, we provide a faster connected components algorithm that is nearly optimal (up to an additive factor of at most $O(\log \log n)$) due to a conditional lower bound of $\Omega(\log d)$ [RVW18, BDE⁺19].

Theorem 3 (Faster Connected Components). *There is an ARBITRARY CRCW PRAM algorithm using $O(m)$ processors that computes the connected components of any given graph. With probability $1 - 1/\text{poly}((m \log n)/n)$, it runs in $O(\log d + \log \log_{m/n} n)$ time.*

For a dense graph with $m = n^{1+\Omega(1)}$, the algorithms in all three theorems run in $O(\log d)$ time with probability $1 - 1/\text{poly}(n)$; if $d = \log_{m/n}^{\Omega(1)} n$, the algorithm in Theorem 3 runs in $O(\log d)$ time.

Note that the time bound in Theorem 1 is dominated by the one in Theorem 3. We include Theorem 1 here in pursuit of simpler proofs of Theorem 2 and Theorem 3.

1.2 Related Work and Technical Overview

In this section, we give an overview of the techniques in our algorithms, highlighting the challenges in the PRAM model and the main novelty in our work compared to that of Andoni et al. [ASS⁺18] and Behnezhad et al. [BDE⁺19].

1.2.1 Related Work

Andoni et al. [ASS⁺18] observed that if every vertex in the graph has a degree of at least $b = m/n$, then one can choose each vertex as a *leader* with probability $\Theta(\log(n)/b)$ to make sure that with high probability, every non-leader vertex has at least 1 leader neighbor. As a result, the number of vertices in the contracted graph is the number of leaders, which is $\tilde{O}(n/b)$ in expectation, leading to double-exponential progress, since we have enough space to make each remaining vertex have degree $\tilde{\Omega}(m/(n/b)) = \tilde{\Omega}(b^2)$ in the next round and $\tilde{\Omega}(b^{2^i})$ after i rounds.³ The process of adding edges is called *expansion*, as it expands the neighbor sets. It can be implemented to run in $O(\log d)$ time. This gives an $O(\log d \log \log_{m/n} n)$ running time.

Behnezhad et al. improved the multiplicative $\log \log_{m/n} n$ factor to an additive factor by streamlining the expansion procedure and double-exponential progress when increasing the space per vertex [BDE⁺19]. Instead of increasing the degree of each vertex *uniformly* to at least b in each round, they allow vertices to have different space budgets, so that the degree lower bound varies on the vertices. They define the *level* $\ell(v)$ of a vertex v to control the *budget* of v (space owned by v): initially each vertex is at level 1 with budget m/n . Levels increase over rounds. Each v is assigned a budget of $b(v) = (m/n)^{c^{\ell(v)}}$ for some fixed constant $c > 1$. The maximal level $L := \log_c \log_{m/n} n$ is such that a vertex at level L must have enough space to find all vertices in its component. Based on this idea, they design an MPC algorithm that maintains the following invariant:

³We use \tilde{O} to hide $\text{polylog}(n)$ factors.

Observation 1.1 ([BDE⁺19]). *With high probability, for any two vertices u and v at distance 2, after 4 rounds, if $\ell(u)$ does not increase then their distance decreases to 1; moreover, the skipped vertex w originally between u and v satisfies $\ell(w) \leq \ell(u)$.*⁴

Behnezhad et al. proved an $O(\log d + \log \log_{m/n} n)$ time bound by a potential-based argument on an arbitrary fixed shortest path P_1 in the original graph: in round 1 put 1 coin on each vertex of P_1 (thus at most $d + 1$ coins in total); for the purpose of analysis only, when inductively constructing P_{i+1} in round i , every skipped vertex on P_i is removed and passes its coins evenly to its successor and predecessor (if exist) on P_i . They claimed that any vertex v still on P_i in round i has at least $1.1^{i-\ell(v)}$ coins based on the following:

Observation 1.2 (Claim 3.12 in [BDE⁺19]). *For any path P_i in round i corresponding to an original shortest path, its first and last vertices are on P_j in round j for any $j > i$; moreover, if a vertex on P_i does not increase level in 4 rounds, then either its predecessor or successor on P_i is skipped.*

The first statement is to maintain the connectivity for every pair of vertices in the original graph, and the second statement is to guarantee that each vertex obtains enough coins in the next round. (Observation 1.2 is seemingly obvious from Observation 1.1, however there is a subtle issue overlooked in [BDE⁺19] that invalidates the statement. Their bound is still valid without changing the algorithm by another potential-based argument, which shall be detailed later in this section.)

Since the maximal level is L , by the above claim, a vertex on such a path with length more than 1 in round $R := 8 \log d + L$ would have at least $1.1^{8 \log d} > d + 1$ coins, a contradiction, giving the desired time bound.

Let us first show a counter-example for Observation 1.2 (not for their time bound). Let the original path P_1 be (v_1, v_2, \dots, v_s) . We want to show that the distance between v_1 and v_s is at most 1 after R rounds, so neither v_1 nor v_s can be skipped during the path constructions over rounds. Suppose v_1 does not increase level in 4 rounds, then for v_1 to obtain enough coins to satisfy the claim, v_2 must be skipped. We call an ordered pair (v_i, v_{i+1}) of consecutive vertices on the path *frozen* if v_i is kept and v_{i+1} is skipped in 4 rounds. So (v_1, v_2) is frozen. Let v_3 (any vertex after v_2 suffices) be the *only* vertex on P_1 directly connecting to v_1 after skipping v_2 in 4 rounds. (Observation 1.1 guarantees that a vertex directly connecting to v_1 must exist but cannot guarantee that there is more than 1 such vertex on P_1 .) Note that v_3 *cannot* be skipped, otherwise v_1 is isolated from P_1 and thus cannot connect to v_s . Now consider two cases. If v_4 is skipped, then pair (v_3, v_4) is frozen. If v_4 is not skipped, then assume v_4 does not increase level in 4 rounds. For v_4 to obtain enough coins, v_5 must be skipped because v_3 cannot be skipped to pass coins to v_4 , so the pair (v_4, v_5) is frozen. Observe that from frozen pair (v_1, v_2) , in either case we get another frozen pair, which propagates inductively to the end of P_1 . If we happen to have a frozen pair (v_{s-1}, v_s) , then v_s must be skipped and isolated from v_1 , a contradiction.

Here is a fix to the above issue (formally stated in Lemma 3.12). Note that only the last vertex v on the path can violate the claim that there are at least $1.1^{i-\ell(v)}$ coins on v in round i . Assuming the claim holds for all vertices on the current path, one can always *drop* (to distinguish from skip) v_s and pass its coins to v_{s-1} (which must be kept) if they are a frozen pair. So v_{s-1} , the new last vertex after 4 rounds, obtains enough coins by the induction hypothesis and the second part of Observation 1.1. By the same argument, the resulting path after R rounds has length at most 1. Observe that we dropped $O(R)$ vertices consecutively located at the end of the path, so the concatenated path connecting v_1 and the original v_s has length $O(R)$. Now applying Observation 1.1 to v_1 , we get that in 4 rounds, either its level increases by 1 or its successor is skipped. Therefore, in $O(R + L)$ rounds, there is no successor of v_1 to be skipped and the graph has diameter at most 1 by a union bound over all the (original) shortest paths.

⁴For simplicity, we ignore the issue of changing the graph and corresponding vertices for now.

1.2.2 Our Contributions

Now we introduce the new algorithmic ideas in our PRAM algorithm with a matching time bound.

The first challenge comes from processor allocation. To allocate different-sized blocks of processors to vertices in each round, there is actually an existing tool called *approximate compaction*, which maps the k distinguished elements in a length- n array one-to-one to an array of length $O(k)$ with high probability [Ak189]. (The vertices to be assigned blocks are *distinguished* and their names are indices in the old array; after indexing them in the new array, one can assign them predetermined blocks.) However, the current fastest (and work-optimal) approximate compaction algorithm takes $O(\log^* n)$ time, introducing a multiplicative factor [GMV91]. To avoid this, our algorithm first reduces the number of vertices to $n/\text{polylog}(n)$ in $O(\log \log_{m/n} n)$ time, then uses approximate compaction to rename the remaining vertices by an integer in $[n/\text{polylog}(n)]$ in $O(\log^* n)$ time. After this, each cell of the array to be compacted owns $\text{polylog}(n)$ processors, and each subsequent compaction (thus processor allocation) can be done in $O(1)$ time [Goo91].

The second challenge is much more serious: it is required by the union bound that any vertex u must connect to *all* vertices within distance 2 with high probability if u does not increase in level. Behnezhad et al. achieve this goal by an algorithm based on constant-time sorting and prefix sum, which require $\Omega(\log n / \log \log n)$ time on an ARBITRARY CRCW PRAM with $\text{poly}(n)$ processors [BH89]. Our solution is based on hashing: to expand a vertex u , hash all vertices within distance 2 from u to a hash table owned by u ; if there is a collision, increase the level of u .⁵ As a result, we are able to prove the following result, which is stronger than Observation 1.1 as it holds deterministically and uses only 1 round:

Observation 1.3 (formally stated in Lemma 3.20). *For any two vertices u and v at distance 2, if $\ell(u)$ does not increase then their distance decreases to 1 in the next round; moreover, any vertex originally between u and v has level at most $\ell(u)$.*

The idea of increasing the level immediately after seeing a collision gives a much cleaner proof of Observation 1.3, but might be problematic in bounding the total space/number of processors: a vertex with many vertices within distance 2 can incur a collision and level increase very often. We circumvent this issue by increasing the level of each budget- b vertex with probability $\tilde{\Theta}(b^{-\delta})$ before hashing. Then a vertex v with at least b^δ vertices within distance 2 would see a level increase with high probability; if the level does not increase, there should be at most b^δ vertices within distance 2, thus there is a collision when expanding v with probability $1/\text{poly}(b)$. As a result, the probability of level increase is $\tilde{\Theta}(b^{-\delta}) + 1/\text{poly}(b) \leq b^{-c}$ for some constant $c > 0$, and we can assign a budget of $b^{1+\Omega(c)}$ to a vertex with increased level, leading to double-exponential progress. As a result, the total space is $O(m)$ with good probability since the union bound is over all $\text{polylog}(n)$ different levels and rounds, instead of $O(n^2)$ shortest paths.

Suitable combination of the ideas above yields a PRAM algorithm that reduces the diameter of the graph to at most 1 in $O(R) = O(\log d + \log \log_{m/n} n)$ time, with one *flexibility*: the relationship between the level $\ell(v)$ and budget $b(v)$ of vertex v in our algorithm is *not* strictly $b(v) = (m/n)^{c^{\ell(v)}}$ for some fixed constant $c > 1$ as in Behnezhad et al. [BDE⁺19]; instead, we allow vertices with the same level to have *two* different budgets. We show that such flexibility still maintains the key invariant of our algorithm (without influencing the asymptotic space bound): if a vertex is not a root in a tree in the labeled digraph, then its level must be strictly lower than the level of its parent (formally stated in Lemma 3.2).⁶ Using hashing and a proper parent-update method, our algorithm does not need to compute the number of neighbors with a certain level for *each* vertex, which is required in [ASS⁺18, BDE⁺19] and solved by constant-time sorting and prefix sum on an MPC. If this were done by a direct application of (constant-time) *approximate counting* (cf.

⁵Hashing also naturally removes the duplicate neighbors to get the desired space bound - a goal achieved by sorting in [ASS⁺18, BDE⁺19].

⁶Our algorithm adopts the framework of *labeled digraph* (or *parent graph*) for computing and representing components, which is standard in PRAM literatures, see §2.

[Ajt90]) on each vertex, then each round would take $\Omega(k)$ time where k is the maximal degree of any vertex, so our new ideas are essential to obtain the desired time bound.

Finally, we note that while it is straightforward to halt when the graph has diameter at most 1 in the MPC algorithm, it is not correct to halt (nor easy to determine) in this case due to the different nature of our PRAM algorithm. After the diameter reaches $O(1)$, to correctly compute components and halt the algorithm, we borrow an idea from [LT19] to flatten all trees in the labeled digraph in $O(R)$ time, then apply our *slower* connected components algorithm (cf. Theorem 1) to output the correct components in $O(\log \log_{m/n} n)$ time, which is $O(R)$ total running time.

1.3 Roadmap

The rest of this paper is organized as follows. In §2, we introduce our framework of the algorithms and some essential building blocks. Our faster connected components algorithm with running time $O(\log d + \log \log_{m/n} n)$ is presented in §3. We include the analysis of its correctness, running time, number of processors, success probability, and implementation in each subsection of §3 respectively. The full proofs of Theorem 1, Theorem 2, and Theorem 3 are deferred to the appendix.

2 Preliminaries

2.1 Framework

We formulate the problem of computing connected components concurrently as follows: label each vertex v with a unique vertex $v.p$ in its component. Such a labeling gives a constant-time test for whether two vertices v and w are in the same component: they are if and only if $v.p = w.p$. We begin with every vertex self-labeled ($v.p = v$) and successively update labels until there is exactly one label per component.

The labels define a directed graph (*labeled digraph*) with arcs $(v, v.p)$, where $v.p$ is the *parent* of v . We maintain the invariant that the only cycles in the labeled digraph are self-loops (arcs of the form (v, v)). Then this digraph consists of a set of rooted trees, with v a root if and only if $v = v.p$. Some authors call the root of a tree the *leader* of all its vertices. We know of only one algorithm in the literature that creates non-trivial cycles, that of Johnson and Metaxis [JM97]. Acyclicity implies that when the parent of a root v changes, the new parent of v is not in the tree rooted at v (for any order of the concurrent parent changes). We call a tree *flat* if the root is the parent of every vertex in the tree. Some authors call flat trees *stars*.

In our connected components and spanning forest algorithms (cf. §B and §C in the appendix), we maintain the additional invariant that if the parent of a non-root v changes, its new parent is in the same tree as v (for any order of the parent changes). This invariant implies that the partition of vertices among trees changes only by set union; that is, no parent change moves a proper subtree to another tree. We call this property *monotonicity*. Most of the algorithms in the literature that have a correct efficiency analysis are monotone. Liu and Tarjan [LT19] analyze some non-monotone algorithms. In our faster connected components algorithm (cf. §3), only the preprocessing and postprocessing stages are monotone, which means the execution between these two stages can move subtrees between different trees in the labeled digraph.

2.2 Building Blocks

Our algorithms use three standard and one not-so-standard building blocks, which link (sub)trees, flatten trees, alter edges, and add edges, respectively. (Classic PRAM algorithms develop many techniques to make the graph sparser, e.g., in [Gaz91, HZ96, HZ01], not denser by adding edges.)

We treat each edge $\{v, w\}$ as a pair of oppositely directed arcs (v, w) and (w, v) . A *direct link* applies to a graph arc (v, w) such that v is a root and w is not in the tree rooted at v ; it makes w the parent of v . A *parent link* applies to a graph arc (v, w) and makes $w.p$ the parent of v ; note that v and $w.p$ are not necessarily roots. Concurrent direct links maintain monotonicity while concurrent parent links do not. We add additional constraints to prevent the creation of a cycle in both cases. Specifically, in the case of parent links, if a vertex is not a root in a tree in the labeled digraph, then its level must be strictly lower than the level of its parent (formally stated in Lemma 3.2).

Concurrent links can produce trees of arbitrary heights. To reduce the tree heights, we use the *shortcut* operation: for each v do $v.p := v.p.p$. One shortcut roughly halves the heights of all trees; $O(\log n)$ shortcuts make all trees flat. Hirschberg et al. [HCS79] introduced shortcutting in their connected components algorithm; it is closely related to the *compress* step in tree contraction [MR85] and to *path splitting* in disjoint-set union [TvL84].

Our third operation changes graph edges. To *alter* $\{v, w\}$, we replace it by $\{v.p, w.p\}$. Links, shortcuts, and edge alterations suffice to efficiently compute components. Liu and Tarjan [LT19] analyze simple algorithms that use combinations of our first three building blocks.

To obtain a good bound for small-diameter graphs, we need a fourth operation that adds edges. We *expand* a vertex u by adding an edge $\{u, w\}$ for a neighbor v of u and a neighbor w of v . The key idea for implementing expansion is hashing, which is presented below.

Suppose each vertex owns a block of K^2 processors. For each processor in a block, we index it by a pair $(p, q) \in [K] \times [K]$. For each vertex u , we maintain a size- K table $H(u)$. We choose a random hash function $h : [n] \rightarrow [K]$. At the beginning of an expansion, for each graph arc (u, v) , we write vertex v into the $h(v)$ -th cell of $H(u)$. Then we can expand u as follows: each processor (p, q) reads vertex v from the p -th cell of $H(u)$, reads vertex w from the q -th cell of $H(v)$, and writes vertex w into the $h(w)$ -th cell of $H(u)$. For each $w \in H(u)$ after the expansion, $\{u, w\}$ is considered an *added* edge in the graph and is treated the same as any other edge.

The key difference between our hashing-based expansion and that in the MPC algorithms is that a vertex w within distance 2 from u might not be in $H(u)$ after the expansion due to a collision, so crucial to our analysis is the way to handle collisions. All hash functions in this paper are pairwise independent, so each processor doing hashing in each round only needs to read two words, which uses $O(1)$ private memory and time.

3 Faster Connected Components Algorithm

In this section we prove Theorem 3 by presenting a faster algorithm for connected components.

Faster Connected Components algorithm: repeat {EXPAND-MAXLINK} until the graph has diameter at most 1 and all trees are flat; run the connected components algorithm from Theorem 1 on the remaining graph.

Each iteration of the repeat loop is called a *round*. The *break condition* that the graph has diameter at most 1 and all trees are flat is tested at the end of each round.

To simplify the presentation, we make the following assumption, which will be removed without influencing the asymptotic running time, number of processors, and success probability (cf. §D in the appendix).

Assumption 3.1. Let $c = 200$, at the beginning of the first round each vertex has a distinct id in $[2m / \log^c n]$ and owns a space block of size $\max\{m/n, \log^c n\} / \log^2 n$.

The remainder of this section is organized as follows. The detailed method EXPAND-MAXLINK is presented in §3.1. In §3.2 we prove that our algorithm correctly computes the connected components of the input graph when it ends. In §3.3, we implement EXPAND-MAXLINK on an ARBITRARY CRCW PRAM in $O(1)$ time. We prove that the algorithm uses $O(m)$ processors over all rounds in §3.4. Finally, we show that the graph diameter is at most 1 and all trees are flat after $O(\log d + \log \log_{m/n} n)$ rounds in §3.5. After the graph diameter reaches 1, it is easy to apply Theorem 1 to output the connected components of the input graph in $O(\log \log_{m/n} n)$ additional time, giving Theorem 3.

3.1 Algorithmic Framework

In this section, we present the algorithmic framework of EXPAND-MAXLINK, the ingredient of each round of Faster Connected Component algorithm.

EXPAND-MAXLINK uses the following three subroutines, which were introduced as building blocks in §2.2:

ALTER: for each edge $e = \{v, w\}$: replace it by $\{v.p, w.p\}$.
SHORTCUT: for each vertex u : update $u.p$ to $u.p.p$.
MAXLINK: repeat {for each vertex v : let $u := \arg \max_{w \in N(v).p} \ell(w)$, if $\ell(u) > \ell(v)$ then update $v.p$ to u } for 2 iterations.

The fourth building block is to expand each vertex v to try to connect to all vertices within distance 2 from v , which corresponds to Steps (3-5). We give detailed explanations of the key concepts and steps after the algorithm.

EXPAND-MAXLINK:

1. MAXLINK; ALTER.
2. For each root v : increase $\ell(v)$ with probability $10 \log n / b(v)^{0.1}$.
3. For each root v : for each root $w \in N(v)$: if $b(w) = b(v)$ then hash w into $H(v)$.
4. For each root v : if there is a collision in $H(v)$ then mark v as *dormant*. For each vertex v : if there is a dormant vertex in $H(v)$ then mark v as *dormant*.
5. For each root v : for each $w \in H(v)$: for each $u \in H(w)$: hash u into $H(v)$. For each root v : if there is a collision in $H(v)$ then mark v as *dormant*.
6. MAXLINK; SHORTCUT; ALTER.
7. For each root v : if v is dormant and did not increase level in Step (2) then increase $\ell(v)$.
8. For each root v : assign a block of size $b_{\ell(v)}$ to v .

Level and budget. The *level* $\ell(v)$ of a vertex v is a non-negative integer that can either remain the same or increase by one during a round. At the beginning of round 1, each vertex v is at level 1 and owns a block of size $b_1 := \max\{m/n, \log^c n\} / \log^2 n$ by Assumption 3.1. During a round, some roots become non-roots by updating their parents. If a vertex remains a root, its level might increase. A root with level ℓ is assigned

a block of size $b_\ell := b_1^{1.01^{\ell-1}}$ at the end of the round. Given b , a vertex v has *budget* $b(v) := b$ if the current block owned by v has size b . Each block of size b is partitioned into \sqrt{b} indexed tables, each of size \sqrt{b} .

Neighbor set. The edges that define the current graph include: (i) the (altered) original edges corresponding to edge processors, and (ii) the (altered) added edges in the tables over all rounds of all vertices. Any vertex within distance 1 of v (including v) in the current graph is called a *neighbor* of v . For any vertex v , let $N(v)$ be the set of its neighbors. In Step (3) we use the old $N(v)$ when initializing the loop that enumerates $N(v)$. For any vertex set S , define $N(S) := \bigcup_{w \in S} N(w)$, and define $S.p := \{w.p \mid w \in S\}$.

Hashing. At the beginning of a round, one random hash function h is chosen. All neighbor roots of all roots use h to do individual hashing in Step (3). A pairwise independent h suffices, so each processor only reads two words. The hashing in Step (5) uses the same h . For each vertex v , let $H(v)$ be the first table in its block, which will store the added edges incident on v . Step (5) is implemented by storing the old tables for all vertices while hashing new items (copied from the old $H(v)$ and old $H(w)$ in the block of w) into the new table.

3.2 Correctness

In this section, we prove that Faster Connected Components algorithm correctly computes the connected components of the input graph when it ends. First of all, we prove a useful property on levels and roots.

Lemma 3.2. *If a vertex v is a non-root at any step, then during the execution after that step, v is a non-root, $\ell(v)$ cannot change, and $1 \leq \ell(v) < \ell(v.p)$.*

Proof. The proof is by an induction on rounds. The lemma clearly holds at the beginning of the first round by definitions.

In Step (1), each iteration of a MAXLINK can only update the parent to a vertex with higher level, which cannot be itself. In Step (2), level increase only applies to roots. The invariant holds after the MAXLINK in Step (6) for the same reasons as above. In SHORTCUT (cf. Step (6)), each vertex v updates its parent to $v.p.p$, which, by the induction hypothesis, must be a vertex with level higher than v if v is a non-root, thus cannot be v . In Step (7), level increase only applies to roots. All other steps and ALTERS do not change the labeled digraph nor levels, giving the lemma. \square

Based on Lemma 3.2, we have the following key result, which implies the correctness of our algorithm.

Lemma 3.3. *The following conditions hold at the end of each round:*

1. *For any component in the input graph, its vertices are partitioned into trees in the labeled digraph such that each tree belongs to exactly one component.*
2. *A tree does not contain all the vertices in its component if and only if there is an edge between a vertex in this tree and a vertex in another tree.*

The detailed proof of Lemma 3.3 is deferred to the appendix.

Lemma 3.4. *If Faster Connected Components algorithm stops, then it correctly computes the connected components of the input graph.*

Proof. When the repeat loop ends, all trees are flat, which means in the last round, all trees are flat before the ALTER in Step (6) since an ALTER does not change the labeled digraph. After the ALTER, all edges are only incident on roots. By Lemma 3.3, these rooted trees partition the connected components; moreover,

if a root is not the only root in its component then there must be an edge between it and another root in its component, i.e., the connectivity of the graph is maintained. This implies that the connected components algorithm from Theorem 1 will correctly compute the connected components of the input graph (see details in the appendix). \square

3.3 Implementation

In this section, we show how to implement EXPAND-MAXLINK on an ARBITRARY CRCW PRAM such that any of the first $O(\log n)$ rounds runs in constant time with good probability.

Lemma 3.5. *With good probability, each of the first $O(\log n)$ rounds can be implemented to run in $O(1)$ time.*

Proof. The ALTER (cf. Steps (1,6)) applies to all edges in the current graph. Since each edge corresponds to a distinct processor, Step (3) and the ALTER take $O(1)$ time.

Steps (2,4,7) and SHORTCUT take $O(1)$ time as each vertex has a corresponding processor and a collision can be detected using the same hash function to check the same location again: there is a collision in $H(v)$ if a vertex w reads a vertex different from w from the $h(w)$ -th cell of $H(v)$ (then w can write a flag to the processor of v to indicate the collision).

In each of the two iterations of MAXLINK, each vertex v updates its parent to a neighbor parent with the highest level if this level is higher than $\ell(v)$. Since a vertex can increase its level by at most 1 in any round (cf. Steps (2,7)), there are $O(\log n)$ different levels. Let each neighbor of v write its parent with level ℓ to the ℓ -th cell of an array of length $O(\log n)$ in the block of v . By the definitions of level and budget, the block of any vertex in any round has size at least $b_1 = \Omega(\log^3 n)$. Therefore, we can assign a processor to each pair of the cells in this array, such that each non-empty cell can determine whether there is a non-empty cell with a larger index in $O(1)$ time. For any non-empty cell, if there is no non-empty cell with a larger index, it must contain a vertex with the maximum level. As a result, Steps (1,6) take $O(1)$ time.

By Step (3), any $w \in H(v)$ has $b(w) = b(v)$, so each $u \in H(w)$ such that $w \in H(v)$ owns a processor in the block of v since $\sqrt{b(v)} \cdot \sqrt{b(w)} = b(v)$ and any vertex in a table is indexed (by its hash value). Therefore, together with collision detection, Step (5) takes $O(1)$ time.

In Step (8), each vertex is assigned a block. The pool of $\Theta(m)$ processors is partitioned into $\Theta(\log^2 n)$ zones such that the processor allocation in round r for vertices with level ℓ uses the zone indexed by (r, ℓ) , where $r, \ell \in O(\log n)$ in the first $O(\log n)$ rounds. Since there are $\Theta(m)$ processors in total and all the vertex ids are in $[2m/\log^c n]$ with good probability (cf. Assumption 3.1), we can use $\Theta(m/\log n)$ processors for each different level and apply approximate compaction (cf. §1.2.2 and see the appendix) to index each root in $O(1)$ time with high probability such that the indices of vertices with the same level are distinct, then assign each of them a distinct block in the corresponding zone. Therefore, Step (8) takes $O(1)$ time with good probability by a union bound over all $O(\log n)$ levels and rounds.

Finally, we need to implement the break condition in $O(1)$ time, i.e., to determine whether the graph has diameter at most 1 and all trees are flat at the end of each round. The algorithm checks the following 2 conditions in each round: (i) all vertices do not change their parents nor levels in this round, and (ii) for any vertices v, w, u such that $w \in H(v)$, $u \in H(w)$ before Step (5), the $h(u)$ -th cell in $H(v)$ already contains u . Conditions (i) and (ii) can be checked in $O(1)$ time by writing a flag to vertex processor v if they do not hold for v , then let each vertex with a flag write the flag to a fixed processor. If there is no such flag then both Conditions (i) and (ii) hold and the loop breaks. If there is a non-flat tree, some parent must change in the SHORTCUT in Step (6). If all trees are flat, they must be flat before the ALTER in Step (6), then an ALTER moves all edges to the roots. Therefore, if Condition (i) holds, all trees are flat and edges are only incident on roots. Moreover, no level changing means no vertex increase its level in Step (2) and there is no dormant vertex in Step (7). So for each root v , $N(v) = H(v)$ after Step (3) and $N(N(v)) = H(v)$

after Step (5) as there is no collision. By Condition (ii), the table $H(v)$ does not change during Step (5), so $N(v) = N(N(v))$. If there exists root v such that there is another root with distance at least 2 from v , then there must exist a vertex $w \neq v$ at distance exactly 2 from v , so $w \notin N(v)$ and $w \in N(N(v))$, contradicting with $N(v) = N(N(v))$. Therefore, any root is within distance at most 1 from all other roots in its component and the graph has diameter at most 1.

Since each step runs in $O(1)$ time with good probability, the lemma follows. \square

3.4 Number of Processors

In this section, we show that with good probability, the first $O(\log n)$ rounds use $O(m)$ processors in total.

First of all, observe that in the case that a root v with level ℓ increases its level in Step (2) but becomes a non-root at the end of the round, v is not assigned a block of size $b_{\ell(v)}$ in Step (8). Instead, v owns a block of size $b_{\ell} = b_{\ell(v)-1}$ from the previous round. Since in later rounds a non-root never participates in obtaining more neighbors by maintaining its table in Steps (3-5) (which is the only place that requires a larger block), such *flexibility* in the relationship between level and budget is acceptable.

By the fact that any root v at the end of any round owns a block of size $b_{\ell(v)} = b_1^{1.01^{\ell(v)-1}}$, a non-root can no longer change its level nor budget (cf. Lemma 3.2), and the discussion above, we obtain:

Corollary 3.6. *Any vertex v owns a block of size b at the end of any round where $b_{\ell(v)-1} = b_1^{1.01^{\ell(v)-2}} \leq b \leq b_1^{1.01^{\ell(v)-1}} = b_{\ell(v)}$; if v is a root, then the upper bound on b is tight.*

Secondly, we prove two simple facts about MAXLINK.

Lemma 3.7. *For any vertex v with parent v' and any $w \in N(v)$ before an iteration of MAXLINK, $\ell(w.p) \geq \ell(v')$ after the iteration; furthermore, if $\ell(w.p) > \ell(v)$ before an iteration, then v must be a non-root after the iteration.*

Proof. For any $w \in N(v)$, its parent has level $\max_{u \in N(w)} \ell(u.p) \geq \ell(v')$ after an iteration of MAXLINK. This implies that if $\ell(w.p) > \ell(v)$ before an iteration, then after that $\ell(v.p)$ is at least the level of the old parent of w which is strictly higher than $\ell(v)$, so v must be a non-root. \square

Lemma 3.8. *For any root v with budget b at the beginning of any round, if there is a root $w \in N(v)$ with at least $b^{0.1}$ neighbor roots with budget b after Step (1), then v either increases level in Step (2) or is a non-root at the end of the round with probability $1 - n^{-5}$.*

Proof. Assume v does not increase level in Step (2). Let w be any root in $N(v)$ after Step (1). Since each root $u \in N(w)$ with budget b (thus level at least $\ell(v)$ by Corollary 3.6) increases its level with probability $10 \log n / b^{0.1}$ independently, with probability at least $1 - (1 - 10 \log n / b^{0.1})^{b^{0.1}} \geq 1 - n^{-5}$, at least one u increases level to at least $\ell(v) + 1$ in Step (2). Since $u \in N(N(v))$ after Step (1), by Lemma 3.7, there is a $w' \in N(v)$ such that $\ell(w'.p) \geq \ell(v) + 1$ after the first iteration of MAXLINK in Step (6). Again by Lemma 3.7, this implies that v cannot be a root after the second iteration and the following SHORTCUT. \square

Using the above result, we can prove the following key lemma, leading to the total number of processors.

Lemma 3.9. *For any root v with budget b at the beginning of any round, $b(v)$ is increased to $b^{1.01}$ in this round with probability at most $n^{-5} + b^{-0.05}$.*

Proof. In Step (2), $\ell(v)$ increases with probability $10 \log n / b^{0.1} \leq b^{-0.08}$ when $c \geq 100$, since $b \geq b_1 \geq \log^{c-2} n$. If $\ell(v)$ does increase here then it cannot increase again in Step (7), so we assume this is not the case (and apply a union bound at the end).

If there is a root $w \in N(v)$ with at least $b^{0.1}$ neighbor roots with budget b after Step (1), then v is a root at the end of the round with probability at most n^{-5} by the assumption and Lemma 3.8. So we assume this is not the case.

By the previous assumption we know that at most $b^{0.1}$ vertices are hashed into $H(v)$ in Step (3). By pairwise independency, with probability at most $(b^{0.1})^2/\sqrt{b} = b^{-0.3}$ there is a collision as the table has size \sqrt{b} , which will increase $\ell(v)$ (cf. Steps (4,7)).

Now we assume that there is no collision in $H(v)$ in Step (3), which means $H(v)$ contains all the at most $b^{0.1}$ neighbor roots with budget b . By the same assumption, each such neighbor root w has at most $b^{0.1}$ neighbor roots with budget b , so there is a collision in $H(w)$ in Step (3) with probability at most $(b^{0.1})^2/\sqrt{b} = b^{-0.3}$. By a union bound over all the $|H(v)| \leq b^{0.1}$ such vertices, v is marked as dormant in the second statement of Step (4) (and will increase level in Step (7)) with probability $b^{-0.2}$.

It remains to assume that there is no collision in $H(v)$ nor in any $H(w)$ such that $w \in H(v)$ after Step (4). As each such table contains at most $b^{0.1}$ vertices, in Step (5) there are at most $b^{0.2}$ vertices to be hashed, resulting in a collision in $H(v)$ with probability at most $(b^{0.2})^2/\sqrt{b} = b^{-0.1}$, which increases $\ell(v)$ in Step (7).

Observe that only a root v at the end of the round can increase its budget, and the increased budget must be $b^{1.01}$ since the level can increase by at most 1 during the round and $b = b_{\ell(v)}$ at the beginning of the round by Corollary 3.6. By a union bound over the events in each paragraph, $b(v)$ is increased to $b^{1.01}$ with probability at most $b^{-0.08} + n^{-5} + b^{-0.3} + b^{-0.2} + b^{-0.1} \leq n^{-5} + b^{-0.05}$. \square

Finally, we are ready to prove an upper bound on the number of processors.

Lemma 3.10. *With good probability, the first $O(\log n)$ rounds use $O(m)$ processors in total.*

Proof. Using Lemma 3.9, by a union bound over all $O(n)$ roots, all $O(\log n)$ rounds, and all $O(\log n)$ different budgets (since there are $O(\log n)$ different levels), with probability at least $1 - n^{-3}$, any root v with budget b at the beginning of any round increases its budget to $b^{1.01}$ with probability at most $b^{-0.05}$. We may assume that the $(1 - n^{-3})$ -probability event always holds since a good-probability result follows from a union bound.

The number of processors for (altered) original edges and vertices are clearly $O(m)$ over all rounds (where each vertex processor needs $O(1)$ private memory to store the corresponding parent, vertex id, hash function, level, and budget). Therefore, we only need to bound the number of processors in blocks that are assigned to a vertex in Step (8) in all $O(\log n)$ rounds. (In the appendix, we will show that the overhead in Step (8) is $O(1)$ with high probability.)

For any positive integer ℓ , let n_ℓ be the number of vertices that ever reaches budget b_ℓ during the first $O(\log n)$ rounds. For any vertex v that ever reaches budget b_ℓ , it has exactly one chance to reach budget $b_{\ell+1}$ in a round if v is a root in that round, which happens with probability at most $b_\ell^{-0.05}$. By a union bound over all $O(\log n)$ rounds, v reaches budget $b_{\ell+1}$ with probability at most $O(\log n) \cdot b_\ell^{-0.05} \leq b_\ell^{-0.04}$ when $c \geq 200$, since $b_\ell \geq b_1 \geq \log^{c-2} n$. We obtain $\mathbb{E}[n_{\ell+1} \mid n_\ell] \leq n_\ell \cdot b_\ell^{-0.04}$, thus by $b_{\ell+1} = b_\ell^{1.01}$, it must be:

$$\mathbb{E}[n_{\ell+1} b_{\ell+1} \mid n_\ell] \leq n_\ell \cdot b_\ell^{-0.04} \cdot b_\ell^{1.01} = n_\ell b_\ell \cdot b_\ell^{-0.03}.$$

By Markov's inequality, $n_{\ell+1} b_{\ell+1} \leq n_\ell b_\ell$ with probability at least $1 - b_\ell^{-0.03} \geq 1 - b_1^{-0.03}$. By a union bound over all $\ell \in O(\log n)$, $n_\ell b_\ell \leq n_1 b_1$ for all $\ell \in O(\log n)$ with probability at least $1 - O(\log n) \cdot b_1^{-0.03} \geq 1 - b_1^{-0.01}$, which is $1 - 1/\text{poly}((m \log n)/n)$ by $b_1 = \max\{m/n, \log^c n\}/\log^2 n$ and $c \geq 100$. So the number of new allocated processors for vertices with any budget in any of the first $O(\log n)$ rounds is at most $n_1 b_1$ with good probability.

Recall from Assumption 3.1 and by a direct calculation, $n_1 \cdot b_1 = O(m/\log^2 n)$ with good probability. Therefore, by a union bound over all the $O(\log n)$ different budgets and $O(\log n)$ rounds, with good probability the total number of processors is $O(m)$. \square

3.5 Diameter Reduction

Let $R := O(\log d + \log \log_{m/n} n)$ where the constant hidden in O will be determined later in this section. The goal is to prove that $O(R)$ rounds of EXPAND-MAXLINK suffice to reduce the diameter of the graph to $O(1)$ and flatten all trees with good probability.

In a high level, our algorithm/proof is divided into the following 3 stages/lemmas:

Lemma 3.11. *With good probability, after round R , the diameter of the graph is $O(R)$.*

Lemma 3.12. *With good probability, after round $O(R)$, the diameter of the graph is at most 1.*

Lemma 3.13. *With good probability, after round $O(R)$, the diameter of the graph is at most 1 and all trees are flat.*

3.5.1 Path Construction

To formalize and quantify the effect of reducing the diameter, consider any shortest path P in the input graph, whose length $|P|$ is at most d . Each ALTER (cf. Steps (1,6)) in each round replaces each vertex on P by its parent, resulting in a path P' of the same length as P . (Note that P' might not be a shortest path in the current graph and can contain loops.) We also add edges to the graph for reducing the diameter of the current graph: for any vertices v and w on path P' , if the current graph contains edge (v, w) , then all vertices exclusively between v and w can be removed from P' , which still results in a valid path in the current graph from the first to the last vertex of P' , reducing the length. If all such paths reduce their lengths to at most d' , the diameter of the current graph is at most d' . In the following, consider any fixed shortest path P_1 at the beginning of round 1. Formally, we have the following inductive construction of paths for diameter reduction:⁷

Definition 3.14 (path construction). *Let all vertices on P_1 be active. For any positive integer r , given path P_r with at least 4 active vertices at the beginning of round r , EXPAND-MAXLINK constructs P_{r+1} by the following 7 phases:*

1. *The ALTER in Step (1) replaces each vertex v on P_r by $v' := v.p$ to get path $P_{r,1}$. For any v' on $P_{r,1}$, let \underline{v}' be on P_r such that $\underline{v}'.p = v'$.*
2. *Let the subpath containing all active vertices on $P_{r,1}$ be $P_{r,2}$.*
3. *After Step (5), set i as 1, and repeat the following until $i \geq |P_{r,2}| - 1$: let $v' := P_{r,2}(i)$, if \underline{v}' is a root and does not increase level during round r then: if the current graph contains edge $(v', P_{r,2}(i+2))$ then mark $P_{r,2}(i+1)$ as skipped and set i as $i+2$; else set i as $i+1$.*
4. *For each $j \in [i+1, |P_{r,2}|+1]$, mark $P_{r,2}(j)$ as passive.*
5. *Remove all skipped and passive vertices from $P_{r,2}$ to get path $P_{r,5}$.*
6. *Concatenate $P_{r,5}$ with all passive vertices on $P_{r,1}$ and $P_{r,2}$ to get path $P_{r,6}$.*
7. *The ALTER in Step (6) replaces each vertex v on $P_{r,6}$ by $v.p$ to get path P_{r+1} .*

For any vertex v on P_r that is replaced by v' in Phase (1), if v' is not skipped in Phase (3), then let \bar{v} be the vertex replacing v' in Phase (7), and call \bar{v} the corresponding vertex of v in round $r+1$.

⁷For any $i \in [|P|+1]$, let $P(i)$ be the i -th vertex on P .

Lemma 3.15. *For any non-negative integer r , the P_{r+1} constructed in Definition 3.14 is a valid path in the graph and all passive vertices are consecutive from the successor of the last active vertex to the end of P_{r+1} .*

Proof. The proof is by an induction on r . Initially, P_1 is a valid path by our discussion on ALTER at the beginning of this section: it only replaces edges by new edges in the altered graph; moreover, the second part of the lemma is trivially true as all vertices are active. Assuming P_r is a valid path and all passive vertices are consecutive from the successor of the last active vertex to the end of the path. We show the inductive step by proving the invariant after each of the 7 phases in Definition 3.14. Phase (1) maintains the invariant. In Phase (2), $P_{r,2}$ is a valid path as all active vertices are consecutive at the beginning of $P_{r,1}$ (induction hypothesis). In Phase (3), if a vertex v is skipped, then there is an edge between its predecessor and successor on the path; otherwise there is an edge between v and its successor by the induction hypothesis; all passive vertices are consecutive from the successor of the last non-skipped vertex to the end of $P_{r,2}$ (cf. Phase (4)), so the invariant holds. In Phase (6), since the first passive vertex on $P_{r,2}$ is a successor of the last vertex on $P_{r,5}$ and the last passive vertex on $P_{r,2}$ is a predecessor of the first passive vertex on $P_{r,1}$ (induction hypothesis), the invariant holds. Phase (7) maintains the invariant. Therefore, P_{r+1} is a valid path and all passive vertices are consecutive from the successor of the last active vertex to the end of P_{r+1} . \square

Now we relate the path construction to the diameter of the graph:

Lemma 3.16. *For any positive integer r , the diameter of the graph at the end of round r is $O(\max_{P_r} |P_{r,2}| + r)$.*

Proof. Let P_1 be from s to t . By an induction on the number of ALTERS and Lemma 3.15, the corresponding vertices of s and t are still connected by path P_{r+1} at the end of round r . Note that by Lemma 3.15, P_{r+1} can be partitioned into two parts after Phase (2): subpath $P_{r,2}$ and the subpath containing only passive vertices. Since in each round we mark at most 2 new passive vertices (cf. Phases (3,4)), we get $|P_{r+1}| \leq |P_{r,5}| + 2r \leq |P_{r,2}| + 2r$. If any path P_{r+1} that corresponds to a shortest path in the original graph have length at most d' , the graph at the end of round r must have diameter at most d' , so the lemma follows. \square

It remains to bound the length of any $P_{r,2}$ in any round r , which relies on the following potential function:

Definition 3.17. *For any vertex v on P_1 , define its potential $\phi_1(v) := 1$. For any positive integer r , given path P_r with at least 4 active vertices at the beginning of round r and the potentials of vertices on P_r , define the potential of each vertex on P_{r+1} based on Definition 3.14 as follows:*

- For each v replaced by $v.p$ in Phase (1), $\phi_{r,1}(v.p) := \phi_r(v)$.
- After Phase (4), for each active vertex v on $P_{r,2}$, if the successor w of v is skipped or passive, then $\phi_{r,4}(v) := \phi_{r,1}(v) + \phi_{r,1}(w)$.
- After Phase (6), for each vertex v on $P_{r,6}$, if v is active on $P_{r,2}$, then $\phi_{r,6}(v) := \phi_{r,4}(v)$, otherwise $\phi_{r,6}(v) := \phi_{r,1}(v)$.
- For each v replaced by $v.p$ in Phase (7), $\phi_{r+1}(v.p) := \phi_{r,6}(v)$.

We conclude this section by some useful properties of potentials.

Lemma 3.18. *For any path P_r at the beginning of round $r \geq 1$, the following holds: (i) $\sum_{v \in P_r} \phi_r(v) \leq d + 1$; (ii) for any v on P_r , $\phi_r(v) \geq 1$; (iii) for any non-skipped v on $P_{r,2}$ and its corresponding vertex \bar{v} on P_{r+1} , $\phi_{r+1}(\bar{v}) \geq \phi_r(v)$.*

Proof. The proof is by an induction on r . The base case follows from $\phi(v) = 1$ for each v on P_1 (cf. Definition 3.17) and $|P_1| \leq d$. For the inductive step, note that by Definition 3.17, the potential of a corresponding vertex is at least the potential of the corresponding vertex in the previous round (and can be larger in the case that its successor is skipped or passive). This gives (ii) and (iii) of the lemma. For any vertex u on $P_{r,2}$, if both u and its successor are active, then $\phi_r(u)$ is presented for exactly 1 time in $\sum_{v \in P_r} \phi_r(v)$ and $\sum_{v \in P_{r+1}} \phi_{r+1}(v)$ respectively; if u is active but its successor w is skipped or passive, then $\phi_r(u) + \phi_r(w)$ is presented for exactly 1 time in each summations as well; if u and its predecessor are both passive, then $\phi_r(u)$ is presented only in $\sum_{v \in P_r} \phi_r(v)$; the potential of the last vertex on $P_{r,2}$ might not be presented in $\sum_{v \in P_{r+1}} \phi_{r+1}(v)$ depending on i after Phase (3). Therefore, $\sum_{v \in P_{r+1}} \phi_{r+1}(v) \leq \sum_{v \in P_r} \phi_r(v)$ and the lemma holds. \square

3.5.2 Proof of Lemma 3.11

First of all, we need an upper bound on the maximal possible level:

Lemma 3.19. *With good probability, the level of any vertex in any of the first $O(\log n)$ rounds is at most $L := 1000 \max\{2, \log \log_{m/n} n\}$.*

Proof. By Lemma 3.10, with good probability the total number of processors used in the first $O(\log n)$ rounds is $O(m)$. We shall condition on this happening then assume for contradiction that there is a vertex v with level at least L in some round.

If $\log \log_{m/n} n \leq 2$, then $m/n \geq n^{1/4}$. By Corollary 3.6, a block owned by v has size at least

$$b_1^{1.01^{2000-2}} \geq b_1^{20} \geq (m/n / \log^2 n)^{20} \geq (n^{1/5})^{20} = n^4 \geq m^2,$$

which is a contradiction as the size of this block owned by v exceeds the total number of processors $O(m)$.

Else if $\log \log_{m/n} n > 2$, then by Corollary 3.6, a block owned by v has size at least

$$b_1^{1.01^{L-2}} \geq b_1^{1.01^{999 \log \log_{m/n} n}} \geq b_1^{(\log_{m/n} n)^{10}} \geq b_1^{8 \log_{m/n} n}. \quad (1)$$

Whether $m/n \leq \log^c n$ or not, if $c \geq 10$, it must be $b_1 = \max\{m/n, \log^c n\} / \log^2 n \geq \sqrt{m/n}$. So the value of (1) is at least $n^4 \geq m^2$, contradiction. Therefore, the level of any vertex is at most L . \square

We also require the following key lemma:

Lemma 3.20. *For any root v and any $u \in N(N(v))$ at the beginning of any round, let u' be the parent of u after Step (1). If v does not increase level and is a root during this round, then $u' \in H(v)$ after Step (5).*

To prove Lemma 3.20, we use another crucial property of the algorithm, which is exactly the reason behind the design of MAXLINK.

Lemma 3.21. *For any root v and any $u \in N(N(v))$ at the beginning of any round, if v does not increase level in Step (2) and is a root at the end of the round, then $u.p$ is a root with budget $b(v)$ after Step (1).*

Proof. By Lemma 3.2, v is a root during this round. For any $w \in N(v)$ and any $u \in N(w)$, applying Lemma 3.7 for 2 times, we get that $\ell(v) \leq \ell(w.p)$ and $\ell(v) \leq \ell(u.p)$ after the MAXLINK in Step (1). If there is a $u \in N(N(v))$ such that $u.p$ is a non-root or $\ell(u.p) > \ell(v)$ before the ALTER in Step (1), it must be $\ell(u.p.p) > \ell(v)$ by Lemma 3.2. Note that $u.p$ is in $N(N(v))$ after the ALTER, which still holds before Step (6) as we only add edges. By Lemma 3.7, there is a $w' \in N(v)$ such that $\ell(w'.p) > \ell(v)$ after the first iteration of MAXLINK in Step (6). Again by Lemma 3.7, this implies that v cannot be a root after the second iteration, a contradiction. Therefore, for any $u \in N(N(v))$, $u.p$ is a root with level $\ell(v)$ (thus budget $b(v)$) after Step (1). \square

With the help of Lemma 3.21 we can prove Lemma 3.20:

Proof of Lemma 3.20. For any vertex u , let $N'(u)$ be the set of neighbors after Step (1). First of all, we show that after Step (5), $H(v)$ contains all roots in $N'(N'(v))$ with budget b , where b is the budget of v at the beginning of the round. For any root $w \in N'(v)$, in Step (3), all roots with budget $b(w)$ in $N'(w)$ are hashed into $H(w)$. If there is a collision in any $H(w)$, then v must be dormant (cf. Step (4)) thus increases level in either Step (2) or (7), contradiction. So there is no collision in $H(w)$ for any $w \in N'(v)$, which means $H(w) \supseteq N'(w)$. Recall that $v \in N'(v)$ and we get that all roots with budget $b(w) = b$ from $N'(N'(v))$ are hashed into $H(v)$ in Step (5). Again, if there is a collision, then v must be dormant and increase level in this round. Therefore, $N'(N'(v)) \subseteq H(v)$ at the end of Step (5).

By Lemma 3.21, for any $u \in N(N(v))$ at the beginning of any round, $u' = u.p$ is a root with budget b in $N'(N'(v))$ after Step (1). Therefore, $u' \in H(v)$ at the end of Step (5), giving Lemma 3.20. \square

The proof of Lemma 3.11 relies on the following lemma based on potentials:

Lemma 3.22. *At the beginning of any round $r \geq 1$, for any active vertex v on any path P_r , $\phi_r(v) \geq 2^{r-\ell(v)}$.*

Proof. The proof is by an induction on r . The base case holds because for any (active) vertex v on P_1 , $\phi_1(r) = 1$ and $r = \ell(v) = 1$. Now we prove the inductive step from r to $r + 1$ given that the corresponding vertex \bar{v} of $v \in P_r$ is on P_{r+1} and active.

Suppose v is a non-root at the end of round r . If v is a non-root at the end of Step (1), then $\ell(v.p) > \ell(v)$ after Step (1) by Lemma 3.2, and $\ell(\bar{v}) \geq \ell(v.p) > \ell(v)$; else if v first becomes a non-root in Step (6), then $\bar{v} = v.p$ and $\ell(v.p) > \ell(v)$ after Step (6) by Lemma 3.2. So by the induction hypothesis, $\phi_{r+1}(\bar{v}) \geq \phi_r(v) \geq 2^{r-\ell(v)} \geq 2^{r+1-\ell(\bar{v})}$.

Suppose v increases its level in round r . Let ℓ be the level of v at the beginning of round r . If the increase happens in Step (2), then v is a root after Step (1). Whether v changes its parent in Step (6) or not, the level of $\bar{v} = v.p$ is at least $\ell + 1$. Else if the increase happens in Step (7), then v is a root after Step (6). So $\bar{v} = v$ and its level is at least $\ell + 1$ at the end of the round. By the induction hypothesis, $\phi_{r+1}(\bar{v}) \geq \phi_r(v) \geq 2^{r-\ell} \geq 2^{r+1-\ell(\bar{v})}$.

It remains to assume that v is a root and does not increase level during round r . By Lemma 3.20, for any $u \in N(N(v))$ at the beginning of round r , the parent u' of u after Step (1) is in $H(v)$ after Step (5). Since v is a root during the round, it remains on $P_{r,2}$ after Phase (2). We discuss two cases depending on whether v is at position before $|P_{r,2}| - 1$ or not.

In Phase (3), note that if $v = P_{r,2}(i)$ where $i < |P_{r,2}| - 1$, then $P_{r,2}(i + 2)$ is the parent of a vertex in $N(N(v))$ after Step (1), which must be in $H(v)$ after Step (5). Therefore, the graph contains edge $(v, P_{r,2}(i + 2))$ and $v' := P_{r,2}(i + 1)$ is skipped, thus $\phi_{r,4}(v) = \phi_{r,1}(v) + \phi_{r,1}(v')$ by Definition 3.17. Since $i + 1 \leq |P_{r,2}| + 1$, v' is an active vertex on $P_{r,1}$. By the induction hypothesis, $\phi_{r,1}(v') = \phi_r(v') \geq 2^{r-\ell(v')}$ (recall that \underline{v}' is replaced by its parent v' in Phase (1)/Step (1)). If $\ell(\underline{v}') > \ell(v)$, then applying Lemma 3.7 for 2 times we get that v is a non-root after Step (1), a contraction. Therefore, $\phi_{r,1}(v') \geq 2^{r-\ell(\underline{v}')} \geq 2^{r-\ell(v)}$ and $\phi_{r,4}(v) \geq \phi_{r,1}(v) + \phi_{r,1}(v') \geq \phi_r(v) + 2^{r-\ell(v)} \geq 2^{r+1-\ell(v)}$.

On the other hand, if $i \geq |P_{r,2}| - 1$ is reached after Phase (3), it must be $i < |P_{r,2}| + 1$ by the break condition of the loop in Phase (3). Note that $v' := P_{r,2}(i + 1)$ is marked as passive in Phase (4), and by Definition 3.17, $\phi_{r,4} = \phi_{r,1}(v) + \phi_{r,1}(v')$. Moreover, since $i + 1 \leq |P_{r,2}| + 1$, v' is an active vertex on $P_{r,1}$. Using the same argument in the previous paragraph, we obtain $\phi_{r,4}(v) \geq 2^{r+1-\ell(v)}$.

By Definition 3.17, after Phase (7), $\phi_{r+1}(\bar{v}) = \phi_{r,6}(v) = \phi_{r,4}(v) \geq 2^{r+1-\ell(v)} = 2^{r+1-\ell(\bar{v})}$. As a result, the lemma holds for any active vertex \bar{v} on P_{r+1} , finishing the induction and giving the lemma. \square

Proof of Lemma 3.11. Let $R := \log d + L$, where L is defined in Lemma 3.19. By Lemma 3.19, with good probability, $\ell(v) \leq L$ for any vertex v in any of the first $O(\log n)$ rounds, and we shall condition on this

happening. By Lemma 3.22, at the beginning of round R , if there is a path P_R of at least 4 active vertices, then for any of these vertices v , it must be $\phi_R(v) \geq 2^{R-\ell(v)} \geq 2^{R-L} \geq d$. So $\sum_{v \in P_R} \phi_r(v) \geq 4d > d + 1$, contradicting with Lemma 3.18. Thus, any path P_R has at most 3 active vertices, which means $|P_{R,2}| \leq 3$ by Definition 3.14. Therefore, by Lemma 3.16, the diameter of the graph at the end of round R is $O(R)$ with good probability. \square

3.5.3 Proof of Lemma 3.13

Based on the graph and any P_R at the beginning of round $R+1$, we need a (much simpler) path construction:

Definition 3.23. For any integer $r > R$, given path P_r with $|P_r| \geq 3$ at the beginning of round r , EXPAND-MAXLINK constructs P_{r+1} by the following:

1. The ALTER in Step (1) replaces each vertex v on P_r by $v' := v.p$ to get path $P_{r,1}$. For any v' on $P_{r,1}$, let \underline{v}' be on P_r such that $\underline{v}'.p = v'$.
2. After Step (5), let $v' := P_{r,1}(1)$, if \underline{v}' is a root at the end of round r and does not increase level during round r then: if the current graph contains edge $(v', P_{r,1}(3))$ then remove $P_{r,1}(2)$ to get path $P_{r,2}$.
3. The ALTER in Step (6) replaces each vertex v on $P_{r,2}$ by $v.p$ to get path P_{r+1} .

For any vertex v on P_r that is replaced by v' in the first step, if v' is not removed in the second step, then let \bar{v} be the vertex replacing v' in the third step, and call \bar{v} the corresponding vertex of v in round $r + 1$.

An analog of Lemma 3.15 immediately shows that P_r is a valid path for any $r \geq R + 1$. The proof of Lemma 3.12 is simple enough without potential:

Proof of Lemma 3.12. By Lemma 3.11, at the beginning of round $R + 1$, with good probability, any P_{R+1} has length $O(R)$. We shall condition on this happening and apply a union bound at the end of the proof. In any round $r > R$, for any path P_r with $|P_r| \geq 3$, consider the first vertex v' on $P_{r,1}$ (cf. Definition 3.23). If \underline{v}' is a non-root or increases its level during round r , then by the first 3 paragraphs in the proof of Lemma 3.22, it must be $\ell(\bar{v}') \geq \ell(\underline{v}') + 1$. Otherwise, by Lemma 3.20, there is an edge between v' and the successor of its successor in the graph after Step (5), which means the successor of v' on $P_{r,1}$ is removed in the second step of Definition 3.23. Therefore, the number of vertices on P_{r+1} is one less than P_r if $\ell(\bar{v}') = \ell(\underline{v}')$ as the level of a corresponding vertex cannot be lower. By Lemma 3.19, with good probability, the level of any vertex in any of the $O(\log n)$ rounds cannot be higher than L . As P_{R+1} has $O(R)$ vertices, in round $r = O(R) + L + R = O(R) \leq O(\log n)$, the number of vertices on any P_r is at most 2. Therefore, the diameter of the graph after $O(R)$ rounds is at most 1 with good probability. \square

Proof of Lemma 3.13. Now we show that after the diameter reaches 1, if the loop has not ended, then the loop must break in $2L + \log_{5/4} L$ rounds with good probability, i.e., the graph has diameter at most 1 and all trees are flat.

For any component, let u be a vertex in it with the maximal level and consider any (labeled) tree of this component. For any vertex v in this tree that is incident with an edge, since the diameter is at most 1, v must have an edge with u , which must be a root. So v updates its parent to a root with the maximal level after a MAXLINK, then any root must have the maximal level in its component since a root with a non-maximal level before the MAXLINK must have an edge to another tree (cf. Lemma 3.3). Moreover, if v is a root, this can increase the maximal height among all trees in its component by 1.

Consider the tree with maximal height ξ in the labeled digraph after Step (1). By Lemmas 3.2 and 3.19, with good probability $\xi \leq L$. The maximal level can increase by at most 1 in this round. If it is increased in Step (7), the maximal height is at most $\lceil \xi/2 \rceil + 1$ after the MAXLINK in the next round; otherwise, the

maximal height is at most $\lceil(\xi + 1)/2\rceil \leq \lceil\xi/2\rceil + 1$. If $\xi \geq 4$, then the maximal height of any tree after Step (1) in the next round is at most $(4/5)\xi$ (the worst case is that a tree with height 5 gets shortcutted to height 3 in Step (6) and increases its height by 1 in the MAXLINK of Step (1) in the next round). Therefore, after $\log_{5/4} L$ rounds, the maximal height of any tree is at most 3.

Beyond this point, if any tree has height 1 after Step (1), then it must have height 1 at the end of the previous round since there is no incident edge on leaves after the ALTER in the previous round, thus the loop must have been ended by the break condition. Therefore, the maximal-height tree (with height 3 or 2) cannot increase its height beyond this point. Suppose there is a tree with height 3, then if the maximal level of vertices in this component does not change during the round, this tree cannot increase its height in the MAXLINK of Step (6) nor that of Step (1) in the next round, which means it has height at most 2 as we do a SHORTCUT in Step (6). So after L rounds, all trees have heights at most 2 after Step (1). After that, similarly, if the maximal level does not increase, all trees must be flat. Therefore, after L additional rounds, all trees are flat after Step (1). By the same argument, the loop must have been ended in the previous round. The lemmas follows immediately from $L = O(R)$ and Lemma 3.12. \square

By Lemma 3.5 and Lemma 3.13, with good probability, in $O(\log d + \log \log_{m/n} n)$ time the graph diameter is at most 1 and the parent of any vertex is a root. Therefore, applying Theorem 1 to the remaining graph of roots and by Lemma 3.4 we proved Theorem 3.

A An $O(\log d \log \log_{m/n} n)$ Bound: Overview and Related Work

A.1 Overview of Andoni et al.

A natural algorithmic framework for connected components is *leader contraction*: randomly choose some vertices as leaders, then for each non-leader, contract it to an arbitrary leader neighbor (if exists); repeat this leader selection and contraction procedure for the contracted graph until each original component contracts to a single leader. Consider the classic leader contraction algorithm proposed by Reif [Rei84]. In each round, each vertex is chosen as a leader with probability $1/2$, thus in expectation at least $1/4$ fraction of the vertices are contracted, leading to an $O(\log n)$ -time algorithm. Andoni et al. [ASS⁺18] observed that if every vertex in the graph has a degree of at least $b = m/n$, then one can choose each vertex as a leader with probability $\Theta(\log(n)/b)$ to make sure that with high probability, every non-leader vertex has at least 1 leader neighbor. As a result, the number of vertices in the contracted graph is the number of leaders which is $\tilde{O}(n/b)$ in expectation, leading to a double-exponential progress as we have enough space to make each remaining vertex to have degree $\tilde{\Omega}(m/(n/b)) = \tilde{\Omega}(b^2)$ in the next round and $\tilde{\Omega}(b^{2^i})$ after i rounds.

Despite a simple and elegant idea in a high level, their detailed algorithm is complicated, even with the strong primitives on an MPC as blackboxes. In each round of making the graph denser (called *expansion* as it expands the neighbor sets), each low-degree vertex needs to broadcast all of its neighbors to all of its low-degree neighbors and an edge is created between two vertices if one vertex is received by another vertex. To implement such broadcasting step, [ASS⁺18] proposed a method based on duplicating and indexing. In particular, for each vertex v , if a low-degree vertex v has degree t' and t low-degree neighbors, then their method makes each neighbor of v t' copies and each low-degree neighbor of v t' duplicates. By sorting, they can index each neighbor of v and each low-degree neighbor of v . If u is the i -th neighbor of v and u' is the j -th low-degree neighbor of v , then match the i -th duplicate of u' and the j -th copy of u . Such matching method can create all edges required by the broadcasting procedure. The main obstacle to implement it on a PRAM is that their indexing and matching steps highly depend on the sorting primitive which though takes constant rounds in the MPC model, requires $\Omega(\log n / \log \log n)$ time on a CRCW PRAM with $\text{poly}(n)$ processors [BH89].

A.2 Our Contribution in Connected Components

One of our main technical contributions is a new expansion procedure that makes *almost all* vertices to have *approximately high* degree. To be more precise, we allow that:

Observation A.1 (formally stated in Lemma B.13). *A small fraction of vertices do not expand their neighbor sets at all, and some vertices stop their expansions after getting moderately high degrees.*

At first glance, this seems to be a natural approach to potentially solve the problem in $O(\log \log_{m/n} n)$ rounds (each round takes $O(\log d)$ time): if one can make an $(1 - b^{-\delta})$ fraction of vertices to have degrees of at least b^δ (instead of b allowed by space), then sampling leaders with probability $\Theta(\log n / b^\delta)$ guarantees that each high-degree vertex to have a leader neighbor with high probability and the expected number of remaining vertices (which include low-degree vertices and leaders) is $\tilde{O}(n/b^\delta)$; this offers each remaining vertex $\tilde{\Omega}(b^{1+\delta})$ space in the next round, and $\tilde{\Omega}(b^{(1+\delta)^i})$ after i rounds, i.e., double-exponential progress. However, there is a major difficulty in obtaining even only n^δ neighbors for almost all vertices. Consider a vertex w with a low-degree neighbor x and a high-degree neighbor y . For some vertex v in the component of w to obtain enough neighbors in the expansion, it is necessary for w 's neighbor set to include y 's neighbors instead of x 's. But the *failure* that w not expanding at all or expanding x instead of y can propagate through the path passing through w , making a potentially large fraction of vertices not obtaining enough neighbors.

We settle this issue by a *distance doubling argument*: by carefully choosing a path from v to w , one can show that some vertices on this path still obtain $b^{\Omega(\delta)}$ neighbors, though these neighbors are not necessarily

y 's; such large neighbor set passes through this path until reaching v in $O(\log d)$ time. In this case, the number of remaining vertices is $n/b^{\Omega(\delta)}$, offering each of them $b^{1+\Omega(\delta)}$ space in the next round, which reaches $\Omega(n)$ after $\log_{1+\Omega(\delta)} \log_b n$ rounds, completing the execution.

The above idea naturally leads to our hashing-based expansion procedure. To assign more space to the remaining vertices, simply hash them to an array so that a vertex at index i gets the i -th predetermined block of space. If the array length is slightly larger, the expected number of vertices that exclusively own a block is a large fraction. Note that a perfect hashing requires $\Omega(\log^* n)$ time [GMW90] - a cost we do not want to pay in each round.⁸ Each step of expansion also uses hashing, which naturally removes the duplicate neighbors to get desired space bound - a goal achieved also by sorting in [ASS⁺18]. We carefully handle the collisions such that each vertex provably reaches moderately high degree with good probability by the distance doubling argument.

A.3 Our Contribution in Spanning Forest

Many existing connected components algorithm can be directly transformed into a spanning forest algorithm. For example in Reif's algorithm [Rei84], one can output the edges corresponding to leader contraction in each step to the spanning forest. However this is not the case here as we also add edges to the graph. The solution in [ASS⁺18] uses several complicated ideas including merging local shortest path trees, which heavily relies on computing the minimum function in constant time - a goal easily achieved by sorting on an MPC but requiring $\Omega(\log \log n)$ time on a PRAM [FMRW85].

We show how to modify our connected components algorithm with an extended expansion procedure to output a spanning forest. Observe that in our previous expansion procedure, if a vertex u does not stop expansion in step i , then the space corresponding to u contains all the vertices within distance 2^i from u . Based on this, we are able to maintain the distance from u to the closest leader in $O(\log d)$ time by the distance doubling argument used before. After determining the distance of each vertex to its closest leader, for each edge (u, v) , if u has distance x to the closest leader, and v has distance $x - 1$ to the closest leader, then we can set v as the parent of u and add edge (u, v) to the spanning forest. (If there are multiple choices of v , we can choose arbitrary one.) Since this parents assignment does not induce any cycle, we can find a subforest of the graph. If we contract all vertices in each tree of such subforest to the unique leader in that tree (which also takes $O(\log d)$ time by shortcutting as any shortest path tree has height at most d), the problem reduces to finding a spanning forest of the contracted graph. Similar to the analysis of our connected components algorithm, we need $O(\log d)$ time to find a subforest in the contracted graph and the number of contraction rounds is $O(\log \log_{m/n} n)$. Thus, the total running time is asymptotically the same as our connected components algorithm.

Our hashing-based approach simplifies the connected components and spanning forest algorithms of Andoni et al., making them potentially practical: despite sorting has a theoretical constant communication round complexity, it can take non-constant rounds (scale with the problem size) and reasonably high running time (counting in the local computations) when implemented in a real distributed system [ABSS15, AS17]. Our approach completely avoids sorting thus should be preferable in practice.

B An $O(\log d \log \log_{m/n} n)$ -time Connected Components Algorithm

In this section we present our connected components algorithm. For simplicity in presentation, we also consider the COMBINING CRCW PRAM [Ak189], whose computational power is between the ARBITRARY CRCW PRAM and MPC. This model is the same as the ARBITRARY CRCW PRAM, except that if several

⁸The idea of hashing to blocks prevents the use of approximate compaction in processor allocation, which we believe is an added bonus to the simplicity and practical efficiency.

processors write to the same memory cell at the same time, the resulting value is a specified symmetric function (e.g., the sum or min) of the individually written values.

We begin with a simple randomized algorithm proposed by Reif [Rei84] but adapted in our framework, which is called Vanilla algorithm. Our main algorithm will use the same framework with a simple preprocessing method to make the graph denser and an elaborated expansion method to add edges before the direct links. We implement the algorithm on a COMBINING CRCW PRAM and then generalize it to run on an ARBITRARY CRCW PRAM in §B.5.

B.1 Vanilla Algorithm

In each iteration of Vanilla algorithm (see below), some roots of trees are selected to be leaders, which becomes the parents of non-leaders after the LINK. A vertex u is called a *leader* if $u.l = 1$.

Vanilla algorithm: repeat {RANDOM-VOTE; LINK; SHORTCUT; ALTER} until no edge exists other than loops.

RANDOM-VOTE: for each vertex u : set $u.l := 1$ with probability $1/2$, and 0 otherwise.

LINK: for each graph arc (v, w) : if $v.l = 0$ and $w.l = 1$ then update $v.p$ to w .

SHORTCUT: for each vertex u : update $u.p$ to $u.p.p$.

ALTER: for each edge $e = \{v, w\}$: replace it by $\{v.p, w.p\}$.

It is easy to see that Vanilla algorithm uses $O(m)$ processors and can run on an ARBITRARY CRCW PRAM. We call an iteration of the repeat loop in the algorithm a *phase*. Clearly each phase takes $O(1)$ time. We obtain the following results:

Definition B.1. A vertex is ongoing if it is a root but not the only root in its component, otherwise it is finished.

Lemma B.2. At the beginning of each phase, each tree is flat and a vertex is ongoing if and only if it is incident with a non-loop edge.

Proof. The proof is by induction on phases. At the beginning, each vertex is in a single-vertex tree and the edges between trees are in the original graph, so the lemma holds. Suppose it holds for phase $k - 1$. After the LINK in phase k , each tree has height at most two since only non-leader root can update its parent to a leader root. The following SHORTCUT makes the tree flat, then the ALTER moves all the edges to the roots. If a root is not the only root in its component, there must be an edge between it and another vertex not in its tree. \square

Lemma B.3. Given a vertex u , after k phases of Vanilla algorithm, u is ongoing with probability at most $(3/4)^k$.

Proof. We prove the lemma by an induction on k . The lemma is true for $k = 0$. Suppose it is true for $k - 1$. Observe that a non-root can never again be a root. For vertex u to be ongoing after k phases, it must be ongoing after $k - 1$ phases. By the induction hypothesis this is true with probability at most $(3/4)^{k-1}$. Furthermore, by Lemma B.2, if this is true, there must be an edge $\{u, v\}$ such that v is ongoing. With probability $1/4$, $u.l = 0$ and $v.l = 1$, then u is finished after phase k . It follows that the probability that, after phase k , u is still ongoing is at most $(3/4)^k$. \square

By Lemma B.3, the following corollary is immediate by linearity of expectation and Markov's inequality:

Corollary B.4. *After k phases of Vanilla algorithm, the number of ongoing vertices is at most $(7/8)^k n$ with probability at least $1 - (6/7)^k$.*

By Lemma B.2, Corollary B.4, and monotonicity, we have that Vanilla algorithm outputs the connected components in $O(\log n)$ time with high probability.

B.2 Algorithmic Framework

In this section we present the algorithmic framework for our connected components algorithm.

For any vertex v in the current graph, a vertex within distance 1 from v in the current graph (which contains the (altered) original edges and the added edges) is called a *neighbor* of v . The set of neighbors of every vertex is maintained during the algorithm (see the implementation of the EXPAND).

Connected Components algorithm: PREPARE; repeat {EXPAND; VOTE; LINK; SHORTCUT; ALTER} until no edge exists other than loops.

PREPARE: if $m/n \leq \log^c n$ for given constant c then run $c \log_{8/7} \log n$ phases of Vanilla algorithm.

EXPAND: for each ongoing u : expand the neighbor set of u according to some rule.

VOTE: for each ongoing u : set $u.l$ according to some rule in $O(1)$ time.

LINK: for each ongoing v : for each w in the neighbor set of v : if $v.l = 0$ and $w.l = 1$ then update $v.p$ to w .

The SHORTCUT and ALTER are the same as those in Vanilla algorithm. The LINK is also the same in the sense that in our algorithm the graph arc (v, w) is added during the EXPAND in the form of adding w to the neighbor set of v . Therefore, Lemma B.2 also holds for this algorithm.

The details of the EXPAND and VOTE will be presented in §B.3 and §B.4, respectively. We call an iteration of the repeat loop after the PREPARE a *phase*. By Lemma B.2, we can determine whether a vertex is ongoing by checking the existence of non-loop edges incident on it, therefore in each phase, the VOTE, LINK, SHORTCUT, and ALTER take $O(1)$ time.

Let $\delta = m/n'$, where n' is the number of ongoing vertices at the beginning of a phase. Our goal in one phase is to reduce n' by a factor of at least a positive constant power of δ with high probability with respect to δ , so we do a PREPARE before the main loop to obtain a large enough δ with good probability:

Lemma B.5. *After the PREPARE, if $m/n > \log^c n$, then $m/n' \geq \log^c n$; otherwise $m/n' \geq \log^c n$ with probability at least $1 - 1/\log^c n$.*

Proof. The first part is trivial since the PREPARE does nothing. By Corollary B.4, after $c \log_{8/7} \log n$ phases, there are at most $n/\log^c n$ ongoing vertices with probability at least $1 - (6/7)^{c \log_{8/7} \log n} \geq 1 - 1/\log^c n$, and the lemma follows immediately from $m \geq n$. \square

We will be focusing on the EXPAND, VOTE, and LINK, so in each phase it suffices to only consider the induced graph on ongoing vertices with current edges. If no ambiguity, we call this induced graph just the graph, call the current edge just the edge, and call an ongoing vertex just a vertex.

In the following algorithms and analyses, we will use the following assumption for simplicity in the analyses.

Assumption B.6. *The number of ongoing vertices n' is known at the beginning of each phase.*

This holds if running on a COMBINING CRCW PRAM with sum as the combining function to compute n' in $O(1)$ time. Later in §B.5 we will show how to remove Assumption B.6 to implement our algorithms on an ARBITRARY CRCW PRAM.

B.3 The Expansion

In this section, we present the method EXPAND and show that almost all vertices have a large enough neighbor set after the EXPAND with good probability.

B.3.1 Setup

Blocks. We shall use a pool of m processors to do the EXPAND. We divide these into $m/\delta^{2/3}$ indexed blocks, where each block contains $\delta^{2/3}$ indexed processors. Since n' and δ are known at the beginning of each phase (cf. Assumption B.6), if a vertex is assigned to a block, then it is associated with $\delta^{2/3}$ (indexed) processors. We map the n' vertices to the blocks by a random hash function h_B . Each vertex has a probability of being the only vertex mapped to a block, and if this happens then we say this vertex *owns* a block.

Hashing. We use a hash table to implement the neighbor set of each vertex and set the size of the hash table as $\delta^{1/3}$, because we need $\delta^{1/3}$ processors for each cell in the table to do an expansion step (see Step (5a) in the EXPAND). We use a random hash function h_V to hash vertices into the hash tables. Let $H(u)$ be the hash table of vertex u . If no ambiguity, we also use $H(u)$ to denote the set of vertices stored in $H(v)$. If u does not own a block, we think that $H(u) = \emptyset$.

We present the method EXPAND as follows:

EXPAND:

1. Each vertex is either *live* or *dormant* in a step. Mark every vertex as *live* at the beginning.
2. Map the vertices to blocks using h_B . Mark the vertices that do not own a block as *dormant*.
3. For each graph arc (v, w) : if v is live before Step (3) then use h_V to hash v into $H(v)$ and w into $H(v)$, else mark w as *dormant*.
4. For each hashing done in Step (3): if it causes a *collision* (a cell is written by different values) in $H(u)$ then mark u as *dormant*.
5. Repeat the following until there is neither live vertex nor hash table getting a new entry:
 - (a) For each vertex u : for each v in $H(u)$: if v is dormant before Step (5a) in this iteration then mark u as *dormant*, for each w in $H(v)$: use h_V to hash w into $H(u)$.
 - (b) For each hashing done in Step (5a): if it causes a collision in $H(u)$ then mark u as *dormant*.

The first four steps and each iteration of Step (5) in the EXPAND take $O(1)$ time. We call an iteration of the repeat loop in Step (5) a *round*. We say a statement holds before round 0 if it is true before Step (3), it holds in round 0 if it is true after Step (4) and before Step (5), and it holds in round i ($i > 0$) if it is true just after i iterations of the repeat loop in Step (5).

Additional notations. We use $\text{dist}(u, v)$ to denote the *distance* between u and v , which is the length of the shortest path from u to v . We use $B(u, \alpha) = \{v \in V \mid \text{dist}(u, v) \leq \alpha\}$ to represent the set of vertices with distance at most α from u . For any $j \geq 0$ and any vertex u , let $H_j(u)$ be the hash table of u in round j .

Consider a vertex u that is dormant after the EXPAND. We call u *fully dormant* if u is dormant before round 0, i.e., u does not own a block. Otherwise, we call u *half dormant*. For a half dormant u , let $i \geq 0$

be the first round u becomes dormant. For u that is live after the EXPAND, let $i \geq 0$ be the first round that its hash table is the same as the table just before round i . The following lemma shows that in this case, the table of u in round $j < i$ contains exactly the vertices within distance 2^j :

Lemma B.7. *For any vertex u that is not fully dormant, let i be defined above, then it must be that $H_i(u) \subseteq B(u, 2^i)$. Furthermore, for any $j \in [0, i - 1]$, $H_j(u) = B(u, 2^j)$.*

Proof. According to the update rule of $H(u)$, it is easy to show that for any integer $j \geq 0$, $H_j(u) \subseteq B(u, 2^j)$ holds by induction. Now we prove that for any $j \in [0, i - 1]$, $H_j(u) = B(u, 2^j)$. We claim that for any vertex v , if v is not dormant in round j , then $H_j(v) = B(v, 2^j)$. The base case is when $j = 0$. In this case, $H_0(v)$ has no collision, so the claim holds. Suppose the claim holds for $j - 1$, i.e., for any vertex v' which is not dormant in round $j - 1$, it has $H_{j-1}(v') = B(v', 2^{j-1})$. Let v be any vertex which is not dormant in round j . Then since there is no collision, $H_j(v)$ should be $\bigcup_{v' \in H_{j-1}(v)} H_{j-1}(v') = \bigcup_{v' \in B(v, 2^{j-1})} B(v', 2^{j-1}) = B(v, 2^j)$. Thus the claim is true, and it implies that for any $j \in [0, i - 1]$, $H_j(u) = B(u, 2^j)$. \square

Lemma B.8. *The EXPAND takes $O(\log d)$ time.*

Proof. By an induction on phases, any path in the previous phase is replaced by a new path with each vertex on the old path replaced by its parent, so the diameter never increases. Since u either is fully dormant or stops its expansion in round i , the lemma follows immediately from Lemma B.7. \square

B.3.2 Neighbor Set Size Lower Bound

We want to show that the table of u in round i contains enough neighbors, but u becomes dormant in round i possibly due to propagations from another vertex in the table of u that is dormant in round $i - 1$, which does not guarantee the existence of collisions in the table of u (which implies large size of the table with good probability). We overcome this issue by identifying the maximal-radius ball around u with no collision nor fully dormant vertex, whose size serves as a size lower bound of the table in round i .

Definition B.9. *For any vertex u that is dormant after the EXPAND, let r be the minimal integer such that there is no collision nor fully dormant vertex in $B(u, r - 1)$.*

Lemma B.10. *If u is fully dormant then $r = 0$. If u is half dormant then $2^{i-1} < r \leq 2^i$.*

Proof. If u is fully dormant, then $r = 0$ since $B(u, 0) = \{u\}$. Suppose u is half dormant. We prove the lemma by induction on i . The lemma holds for $i = 0$ because $r > 0$ and if $r \geq 2$ then u cannot be dormant in Step (3) nor Step (4). The lemma also holds for $i = 1$ because if $r = 1$ then u becomes dormant in Step (3) or Step (4) and if $r \geq 3$ then u cannot be dormant in round 1.

Suppose $i \geq 2$. Assume $r \leq 2^{i-1}$ and let $v \in B(u, r)$ be a fully dormant vertex or a vertex that causes a collision in $B(u, r)$. Assume u becomes dormant after round $i - 1$. By Lemma B.7, we know that $H_{i-1}(u) = B(u, 2^{i-1})$. Since $r \leq 2^{i-1}$, there is no collision in $B(u, r)$ using h_V . Thus, there is a fully dormant vertex v in $B(u, r) \subseteq H_{i-1}(u)$. Consider the first round $j \leq i - 1$ that v is added into $H(u)$. If $j = 0$, then u is marked as dormant in round 0 by Step (3). If $j > 0$, then in round j , there is a vertex v' in $H_{j-1}(u)$ such that $v \in H_{j-1}(v')$, and v is added into $H_j(u)$ by Step (5a). In this case, u is marked as dormant in round j by Step (5a). In both cases $j = 0$ and $j > 0$, u is marked as dormant in round $j \leq i - 1$ which contradicts with the definition of i . So the only way for u to become dormant in round i is for a vertex v to exist in $H_{i-1}(u)$ which is dormant in round $i - 1$. Assume for contradiction that $r > 2^i$, then by Definition B.9 there is no collision nor fully dormant vertex in $B(u, 2^i)$. By the induction hypothesis, there exists either a collision or a fully dormant vertex in $B(v, 2^{i-1})$. By Lemma B.7, we know that $H_{i-1}(u) = B(u, 2^{i-1})$. It means that $B(v, 2^{i-1}) \subseteq B(u, 2^i)$ contains a collision or a fully dormant vertex, contradiction. \square

To state bounds simply, let $b := \delta^{1/18}$, then hash functions h_B and h_V are from $[n]$ to $[m/b^{12}]$ and from $[n]$ to $[b^6]$, respectively. Note that h_B and h_V need to be independent with each other, but each being pairwise independent suffices, so each processor doing hashing only reads two words.

Lemma B.11. *For any vertex u that is dormant after the EXPAND, $|B(u, r)| \leq b^2$ with probability at most b^{-2} .*

Proof. Let j be the maximal integer such that $j \leq d$ and $|B(u, j)| \leq b^2$. We shall calculate the probability of $r \leq j$, which is equivalent to the event $|B(u, r)| \leq b^2$.

The expectation of the number of collisions in $B(u, j)$ using h_V is at most $\binom{b^2}{2}/b^6 \leq b^{-2}/2$, then by Markov's inequality, the probability of at least one collision existing is at most $b^{-2}/2$.

For any u to be fully dormant, at least one of the $n' - 1$ vertices other than u must have hash value $h_B(u)$. By union bound, the probability of u being fully dormant is at most

$$\frac{n' - 1}{m/b^{12}} \leq \frac{n'}{n'b^6} = \frac{1}{b^6}, \quad (2)$$

where the first inequality follows from $m/n' = b^{18}$. Taking union bound over all vertices in $B(u, j)$ we obtain the probability as at most $b^2 \cdot b^{-6} = b^{-4}$.

Therefore, for any dormant vertex u , $\Pr[|B(u, r)| \leq b^2] = \Pr[r \leq j] \leq b^{-2}/2 + b^{-4} \leq b^{-2}$ by a union bound. \square

Based on Lemma B.11, we shall show that any dormant vertex has large enough neighbor set after the EXPAND with good probability. To prove this, we need the following lemma:

Lemma B.12. $|B(u, r - 1)| \leq |H_i(u)|$.

Proof. Let w be a vertex in $B(u, r - 1)$. If $\text{dist}(w, u) \leq 2^{i-1}$ then $w \in H_{i-1}(u)$ by Lemma B.7 and Definition B.9, and position $h_V(w)$ in $H_i(u)$ remains occupied in round i . Suppose $\text{dist}(w, u) > 2^{i-1}$. Let x be a vertex on the shortest path from u to w and x satisfies $\text{dist}(x, w) = 2^{i-1}$. We obtain $B(x, 2^{i-1}) \subseteq B(u, r - 1)$ since any $y \in B(x, 2^{i-1})$ has $\text{dist}(y, u) \leq \text{dist}(y, x) + \text{dist}(x, u) \leq 2^{i-1} + r - 1 - 2^{i-1}$, and thus $w \in H_{i-1}(x)$ and $x \in H_{i-1}(u)$. So position $h_V(w)$ in $H_i(u)$ remains occupied in round i as a consequence of Step (5a). Therefore Lemma B.12 holds. \square

Lemma B.13. *After the EXPAND, for any dormant vertex u , $|H(u)| < b$ with probability at most b^{-1} .*

Proof. By Lemma B.11, $|B(u, r)| > b^2$ with probability at least $1 - b^{-2}$. If this event happens, u must be half dormant. In the following we shall prove that $|H(u)| \geq b$ with probability at least $1 - b^{-4}$ conditioned on this, then a union bound gives the lemma.

If $i = 0$ then $r = 1$ by Lemma B.10, which gives $|B(u, 1)| > b^2 \geq b$. Consider hashing arbitrary b vertices of $B(u, 1)$. The expectation of the number of collisions among them is at most $b^2/b^6 = b^{-4}$, then by Markov's inequality the probability of at least one collision existing is at most b^{-4} . Thus with probability at least $1 - b^{-4}$ these b vertices all get distinct hash values, which implies $|H(u)| \geq b$. So the lemma holds.

Suppose $i \geq 1$. If $|B(u, r - 1)| \geq b$ then the lemma holds by Lemma B.12. Otherwise, since $|B(u, r)| > b^2$, by Pigeonhole principle there must exist a vertex w at distance $r - 1$ from u such that $|B(w, 1)| > b$. Then by an argument similar to that in the second paragraph of this proof we have that $H_0(w) \geq b$ with probability at least $1 - b^{-4}$. If $i = 1$ then $r = 2$. Since $w \in H_0(u)$ due to Definition B.9, at least b positions will be occupied in $H_1(u)$, and the lemma holds.

Suppose $i \geq 2$. Let w_0 be a vertex on the shortest path from u to w such that $\text{dist}(w_0, w) = 1$. Recursively, for $j \in [1, i - 2]$, let w_j be a vertex on this path such that $\text{dist}(w_{j-1}, w_j) = 2^j$. We have that

$$\text{dist}(w, w_{i-2}) = \text{dist}(w, w_0) + \sum_{j \in [i-2]} \text{dist}(w_{j-1}, w_j) = 2^{i-1} - 1.$$

Thus

$$\text{dist}(w_{i-2}, u) = r - 1 - \text{dist}(w, w_{i-2}) = r - 2^{i-1} \leq 2^{i-1},$$

where the last inequality follows from Lemma B.10. It follows that $w_{i-2} \in H_{i-1}(u)$ by Lemma B.7 and Definition B.9. We also need the following claim, which follows immediately from $\text{dist}(w_j, u) \leq r - 2^j$ and Definition B.9:

Claim B.14. $H_j(w_j) = B(w_j, 2^j)$ for all $j \in [0, i - 2]$.

Now we claim that $|H_{i-1}(w_{i-2})| \geq b$ and prove it by induction, then $|H_i(u)| \geq b$ holds since $w_{i-2} \in H_{i-1}(u)$. $|H_1(w_0)| \geq b$ since $w \in H_0(w_0)$ and $|H_0(w)| \geq b$. Assume $|H_j(w_{j-1})| \geq b$. Since $w_{j-1} \in H_j(w_j)$ by $\text{dist}(w_{j-1}, w_j) = 2^j$ and Claim B.14, we obtain $|H_{j+1}(w_j)| \geq b$. So the induction holds and $|H_{i-1}(w_{i-2})| \geq b$.

By the above paragraph and the first paragraph of this proof we proved Lemma B.13. \square

B.4 The Voting

In this section, we present the method VOTE and show that the number of ongoing vertices decreases by a factor of a positive constant power of b with good probability.

VOTE: for each vertex u : initialize $u.l := 1$,

1. If u is live after the EXPAND then for each vertex v in $H(u)$: if $v < u$ then set $u.l := 0$.
2. Else set $u.l := 0$ with probability $1 - b^{-2/3}$.

There are two cases depending on whether u is live. In Case (1), by Lemma B.7, $H(u)$ must contain all the vertices in the component of u , and so does any vertex in $H(u)$, because otherwise u is dormant. We need to choose the same parent for all the vertices in this component, which is the minimal one in this component as described: a vertex u that is not minimal in its component would have $u.l = 0$ by some vertex v in $H(u)$ smaller than u . Thus all live vertices become finished in the next phase.

In Case (2), u is dormant. Then by Lemma B.13, $|H(u)| \geq b$ with probability at least $1 - b^{-1}$. If this event happens, the probability of no leader in $H(u)$ is at most $(1 - b^{-2/3})^b \leq \exp(-b^{-1/3}) \leq b^{-1}$.

The number of vertices in the next phase is the sum of: (i) the number of dormant leaders, (ii) the number of non-leaders u with $|H(u)| < b$ and no leader in $H(u)$, and (iii) the number of non-leaders u with $|H(u)| \geq b$ and no leader in $H(u)$. We have that the expected number of vertices in the next phase is at most

$$n' \cdot (b^{-2/3} + b^{-1} + (1 - b^{-1}) \cdot b^{-1}) \leq n' \cdot b^{-1/2}.$$

By Markov's inequality, the probability of having more than $n' \cdot b^{-1/4}$ vertices in the next phase is at most

$$\frac{n' \cdot b^{-1/2}}{n' \cdot b^{-1/4}} \leq b^{-1/4}. \quad (3)$$

B.5 Removing the Assumption

In this section, we remove Assumption B.6 which holds on a COMBINING CRCW PRAM, thus generalize the algorithm to run on an ARBITRARY CRCW PRAM.

Recall that we set $b = \delta^{1/18}$ where $\delta = m/n'$ is known. The key observation is that in each phase, all results still hold when we use any \tilde{n} to replace n' as long as $\tilde{n} \geq n'$ and b is large enough. This effectively

means that we use $b = (m/\tilde{n})^{1/18}$ for the hash functions h_B and h_V . This is because the only places that use n' as the number of vertices in a phase are:

1. The probability of a dormant vertex u having $B(u, r) \leq b^2$ (Lemma B.11). Using \tilde{n} to rewrite Inequality (2), we have:

$$\frac{n' - 1}{m/b^{12}} = \frac{n' - 1}{\tilde{n}b^6} \leq \frac{n'}{n'b^6} = \frac{1}{b^6},$$

followed from $m/\tilde{n} = b^{18}$ and $\tilde{n} \geq n'$. Therefore Lemma B.11 still holds.

2. The probability of having more than $n' \cdot b^{-1/4}$ vertices in the next phase. Now we measure the progress by the decreasing in \tilde{n} . The expected number of vertices in the next phase is still at most $n' \cdot b^{-1/4}$, where n' is the exact number of vertices. Therefore by Markov's inequality we can rewrite Inequality (3) as:

$$\frac{n' \cdot b^{-1/2}}{\tilde{n} \cdot b^{-1/4}} \leq b^{-1/4},$$

which is the probability of having more than $\tilde{n} \cdot b^{-1/4}$ vertices in the next phase.

As a conclusion, if $\tilde{n} \geq n'$ and b is large enough in each phase, all analyses still apply. Let c be the value defined in PREPARE. We give the *update rule* of \tilde{n} :

Update rule of \tilde{n} :

If $m/n \leq \log^c n$ then set $\tilde{n} := n/\log^c n$ for the first phase (after the PREPARE), else set $\tilde{n} := n$.
At the beginning of each phase, update $\tilde{n} := \tilde{n}/b^{1/4}$ then update $b := (m/\tilde{n})^{1/18}$.

So $b \geq \log^{c/18} n$ is large enough. By the above argument, we immediately have the following:

Lemma B.15. *Let \tilde{n} and n' be as defined above in each phase. If $\tilde{n} \geq n'$ in a phase, then with probability at least $1 - b^{-1/4}$, $\tilde{n} \geq n'$ in the next phase.*

By an induction on phases, Lemma B.15, and a union bound, the following lemma is immediate:

Lemma B.16. *Given any integer $t \geq 2$, if $\tilde{n} \geq n'$ in the first phase, then $\tilde{n} \geq n'$ in all phases before phase t with probability at least $1 - \sum_{i \in [t-2]} b_i^{-1/4}$, where b_i is the parameter b in phase $i \geq 1$.*

B.6 Running Time

In this section, we compute the running time of our algorithm and the probability of achieving it.

Lemma B.17. *After the PREPARE, if $\tilde{n} \geq n'$ in each phase, then the algorithm outputs the connected components in $O(\log \log_{m/n_1} n)$ phases, where n_1 is the \tilde{n} in the first phase.*

Proof. Let n_i be the \tilde{n} in phase i . By the update rule of \tilde{n} , we have $n_{i+1} \leq n_i/(m/n_i)^{1/72}$, which gives $m/n_{t+1} \geq (m/n_1)^{(73/72)^t}$. If $t = \lceil \log_{73/72} \log_{m/n_1} m \rceil + 1$ then $n_{t+1} < 1$, which leads to $n' = 0$ at the beginning of phase $t+1$. By Lemma B.2 and monotonicity, the algorithm terminates and outputs the correct connected components in this phase since no parent changes. \square

Proof of Theorem 1. We set $c = 100$ in the PREPARE. If $m/n > \log^c n$, then $\tilde{n} = n$ in the first phase by Lemma B.5 and the update rule. Since $\delta = m/n_1 \geq \log^c n$, we have that $b \geq \delta^{1/18} \geq \log^5 n$ in all phases. By Lemma B.16, $\tilde{n} \geq n'$ in all phases before phase $\log n$ with probability at least $1 - \log n \cdot b^{-1/4} =$

$1 - 1/\text{poly}(m \log n/n)$. If this event happens, by Lemma B.17, the number of phases is $O(\log \log_{m/n} n)$. By Lemma B.8, the total running time is $O(\log d \log \log_{m/n} n)$ with good probability.

If $m/n \leq \log^c n$, then by Lemma B.5 and the update rule of \tilde{n} , after the PREPARE which takes time $O(\log \log n)$, with probability at least $1 - 1/\log^c n$ we have $m/n_1 \geq \log^c n$. If this happens, by the argument in the previous paragraph, with probability at least $1 - 1/\text{poly}(m/n_1 \cdot \log n)$, Connected Components algorithm takes time $O(\log d \log \log_{m/n} n)$. Taking a union bound, we obtain that with probability at least $1 - 1/\text{poly}(m/n_1 \cdot \log n) - 1/\log^c n \geq 1 - 1/\text{polylog}(n) = 1 - 1/\text{poly}(m \log n/n)$, the total running time is $O(\log \log n) + O(\log d \cdot \log \log_{m/n_1} n) = O(\log d \log \log_{m/n} n)$. \square

C Spanning Forest Algorithm

In this section, we discuss how to extend our Connected Components algorithm to output a spanning forest, proving Theorem 2.

Since the EXPAND method will add new edges which are not in the input graph, our connected components algorithm cannot give a spanning forest algorithm directly. To output a spanning forest, we only allow direct links on graph arcs of the input graph. However, we want to link many graph arcs concurrently to make a sufficient progress. We extend the EXPAND method to a new subroutine which can link many graph arcs concurrently. Furthermore, after applying the subroutine, there is no cycle induced by link operations, and the tree height is bounded by the diameter of the input graph.

For each arc e in the current graph, we use \hat{e} to denote the original arc in the input graph that is altered to e during the execution. Each edge processor is identified by an original arc \hat{e} and stores the corresponding e in its private memory during the execution. To output the spanning forest, for a original graph arc $\hat{e} = (v, w)$, if at the end of the algorithm, $\hat{e}.f = 1$, then it denotes that the graph edge $\{v, w\}$ is in the spanning forest. Otherwise if both $\hat{e}.f = 0$ and $\hat{e}'.f = 0$ where \hat{e}' denotes a graph arc (w, v) , then the edge $\{v, w\}$ is not in the spanning forest.

C.1 Vanilla Algorithm for Spanning Forest

Firstly, let us see how to extend Vanilla algorithm to output a spanning forest. The extended algorithm is called Vanilla-SF algorithm (see below).

The RANDOM-VOTE and SHORTCUT are the same as those in Vanilla algorithm. In the ALTER we only alter the current edge as in Vanilla algorithm but keep the original edge untouched. We add a method MARK-EDGE and an attribute e for each vertex v to store the current arc that causes the link on v , then $v.\hat{e}$ is the original arc in the input graph corresponding to $v.e$. The LINK is the same except that we additionally mark the original arc in the forest using attribute f if the corresponding current arc causes the link.

Vanilla-SF algorithm: repeat {RANDOM-VOTE; MARK-EDGE; LINK; SHORTCUT; ALTER} until no edge exists other than loops.

MARK-EDGE: for each current graph arc $e = (v, w)$: if $v.l = 0$ and $w.l = 1$ then update $v.e$ to e and update $v.\hat{e}$ to \hat{e} .

LINK: for each ongoing u : if $u.e = (u, w)$ exists then update $u.p$ to w and update $u.\hat{e}.f := 1$.

The digraph defined by the labels is exactly the same as in Vanilla algorithm, therefore Lemma B.2 holds for Vanilla-SF algorithm. It is easy to see that Vanilla-SF algorithm uses $O(m)$ processors and can run on an ARBITRARY CRCW PRAM. Each phase takes $O(1)$ time.

Definition C.1. For any positive integer j , at the beginning of phase j , let F_j be the graph induced by all the edges corresponding to all the arcs \hat{e} with $\hat{e}.f = 1$.

By the execution of the algorithm, for any positive integers $i \leq j$, the set of the edges in F_i is a subset of the set of the edges in F_j .

Lemma C.2. *For any positive integer j , at the beginning of phase j , each vertex u is in the component of $u.p$ in F_j .*

Proof. The proof is by induction. In the first phase, the lemma is vacuously true since $u.p = u$ for all vertices u . Now, suppose the lemma is true at the beginning of phase i . In phase i , if $u.p$ does not change, the claim is obviously true. Otherwise, there are two cases: (i) $u.p$ is changed in the LINK, or (ii) $u.p$ is changed in the SHORTCUT. If $u.p$ is changed in the SHORTCUT, then by Lemma B.2, the original $u.p$ is changed in the LINK, and since u and the original $u.p$ is in the same component of F_i , we only need to show that the LINK does not break the invariant. In the LINK, if $u.p$ is updated to w , there is a current graph arc $u.e = (u, w)$, and thus there exists a graph arc $u.\hat{e} = (x, y)$ in the input graph such that $x.p = u$ and $y.p = w$ at the beginning of phase i . By the induction hypothesis, x and u are in the same component of F_i and thus of F_{i+1} . Similar argument holds for y and w . Since $u.\hat{e}.f$ is set to 1 in the LINK, x and y are in the same component in F_{i+1} . Thus, u and w are in the same component in F_{i+1} . \square

Lemma C.3. *For any positive integer j , F_j is a forest. And in each tree of F_j , there is exactly one root.*⁹

Proof. By the LINK, every time the size of $\{u \mid u.p = u\}$ decreases by 1, the size of $\{\hat{e} \mid \hat{e}.f = 1\}$ increases by 1. Thus the size of $\{\hat{e} \mid \hat{e}.f = 1\}$ is exactly $n - |\{u \mid u.p = u\}|$, which induce at least $|\{u \mid u.p = u\}|$ components in F_j . By Lemma C.2, there are at most $|\{u \mid u.p = u\}|$ components in F_j . So there are exactly $|\{u \mid u.p = u\}|$ components in F_j and each such component must be a tree, and each component contains exactly one vertex u with $u.p = u$. \square

C.2 Algorithmic Framework

In this section, we show how to extend our Connected Components algorithm to a spanning forest algorithm.

Spanning Forest algorithm: FOREST-PREPARE; repeat {EXPAND; VOTE; TREE-LINK; TREE-SHORTCUT; ALTER} until no edge exists other than loops.

FOREST-PREPARE: if $m/n \leq \log^c n$ for given constant c then run Vanilla-SF algorithm for $c \log_{8/7} \log n$ phases.

TREE-LINK: for each ongoing u : update $u.p, u.e, u.\hat{e}$, and $u.\hat{e}.f$ according to some rule.

TREE-SHORTCUT: repeat {SHORTCUT} until no parent changes.

The EXPAND, VOTE, SHORTCUT, and ALTER are the same as in our connected component algorithm.

Similarly to that in Connected Components algorithm, the FOREST-PREPARE makes the number of ongoing vertices small enough with good probability. As analyzed in §B.3 and §B.4, the EXPAND takes $O(\log d)$ time, and VOTE takes $O(1)$ time. TREE-SHORTCUT takes $O(\log d')$ time where d' is the height of the highest tree after the TREE-LINK. Later, we will show that the height d' is $O(d)$, giving $O(\log d)$ running time for each phase.

C.3 The Tree Linking

In this section, we present the method TREE-LINK. The purpose of the TREE-LINK is two-fold: firstly, we want to add some edges to expand the current forest (a subgraph of the input graph); secondly, we need

⁹The *tree* we used here in the forest of the graph should not be confused with the tree in the digraph defined by labels.

the information of these added edges to do link operation of ongoing vertices to reduce the total number of ongoing vertices. For simplicity, all vertices discussed in this section are ongoing vertices in the current phase.

Similar to the EXPAND, we shall use a pool of m processors to do the TREE-LINK. Let n' be the number of vertices. We set all the parameters as the same as in the EXPAND: $\delta = m/n'$, the processors are divided into $m/\delta^{2/3}$ indexed blocks where each block contains $\delta^{2/3}$ indexed processors, and both h_B, h_V are the same hash functions used in the EXPAND. Comparing to Connected Components algorithm, we store not only the final hash table $H(u)$ of u , but also the hash table $H_j(u)$ of u in each round $j \geq 0$. (In Connected Components algorithm, H_j is an analysis tool only.) Let T denote the total number of rounds in Step (5) of the EXPAND.

We present the method TREE-LINK as follows, which maintains: (i) the largest integer $u.\alpha$ for each vertex u such that there is neither collisions, leaders, nor fully dormant vertices in $B(u, u.\alpha)$; (ii) $u.\beta$ as the distance to the nearest leader v (if exists in $B(u, u.\alpha + 1)$) from u ; (iii) a hash table $Q(u)$ to store all vertices in $B(u, u.\alpha)$, which is done by reducing the radius by a factor of two in each iteration and attempting to expand the current $Q(u)$ to a temporary hash table $Q'(u)$.

TREE-LINK:

1. For each vertex u :
 - (a) If $u.l = 1$, set $u.\alpha := -1$ and set hash table $Q(u) := \emptyset$.
 - (b) If $u.l = 0$ and u does not own a block, set $u.\alpha := -1$ and set $Q(u) := \emptyset$.
 - (c) If $u.l = 0$ and u owns a block, set $u.\alpha := 0$ and use h_V to hash u into $Q(u)$.
2. For $j = T \rightarrow 0$: for each vertex u with $u.\alpha \geq 0$: if every v in table $Q(u)$ is live in round j of Step (5) of the EXPAND:
 - (a) Initialize $Q'(u) := \emptyset$.
 - (b) For each v in $Q(u)$: for each w in $H_j(v)$: use h_V to hash w into $Q'(u)$.
 - (c) If there is neither collisions nor leaders in $Q'(u)$ then set $Q(u)$ to be $Q'(u)$ and increase $u.\alpha$ by 2^j .
3. For each current graph arc (v, w) : if $v.l = 1$ then mark w as a *leader-neighbor*.
4. For each vertex u :
 - (a) If $u.l = 1$ then set $u.\beta := 0$.
 - (b) If $u.l = 0$ and $Q(u)$ contains a vertex w marked as a leader-neighbor then set $u.\beta := u.\alpha + 1$.
5. For each current graph arc $e = (v, w)$: if $v.\beta = w.\beta + 1$ then write e into $v.e$ and write \hat{e} to $v.\hat{e}$.
6. For each vertex u : if $u.e = (u, w)$ exists then update $u.p$ to w and update $u.\hat{e}.f := 1$.

For simplicity, if there is no ambiguity, we also use $Q(u)$ to denote the set of vertices which are stored in the table $Q(u)$. We call each iteration j from T down to 0 in Step (2) a *round*.

Lemma C.4. *In any round in the TREE-LINK, for any vertex u , $Q(u) = B(u, u.\alpha)$. Furthermore, at the*

end of the TREE-LINK, $u.\alpha$ is the largest integer such that there is neither collisions nor dormant vertices in $B(u, u.\alpha)$.

Proof. Firstly, we show that $Q(u) = B(u, u.\alpha)$. Our proof is by an induction on j of Step (2). The base case holds since before Step (2), the initialization of $Q(u)$ and $u.\alpha$ satisfy the claim. Now suppose $Q(u) = B(u, u.\alpha)$ at the beginning of round j of Step (2). There are two cases. The first case is that $Q(u)$ does not change in this round. In this case the claim is true by the induction hypothesis. The second case is that $Q(u)$ will be set to $Q'(u)$ in Step (2c). Since there is no collision in $Q'(u)$, we have $Q'(u) = \bigcup_{v \in Q(u)} H_j(v)$, where $Q(u)$ is not set in Step (2c) yet. Since every v in $Q(u)$ is live in round j of Step (5) of the EXPAND, we have $H_j(v) = B(v, 2^j)$ according to Lemma B.7. Together with the induction hypothesis, $Q'(u) = \bigcup_{v \in B(u, u.\alpha)} B(v, 2^j) = B(u, u.\alpha + 2^j)$. Thus by Step (2c), after updating $Q(u)$ and $u.\alpha$, the first part of the lemma holds. Now we prove the second part of the lemma. For convenience in the notation, we say that $B(u, \alpha)$ satisfies property \mathcal{P} if and only if there is neither collisions, leaders, nor fully dormant vertices in $B(u, \alpha)$. We shall show that at the end of the TREE-LINK, $u.\alpha$ is the largest integer such that $B(u, u.\alpha)$ satisfies \mathcal{P} . For any vertex u , if $u.l = 0$ and u is fully dormant, or $u.l = 1$, then the claim holds due to Step (1). Consider a vertex u with $u.l = 0$ and u owning a block. Our proof is by induction on j of Step (2). We claim that at the end of round j of Step (2), the following two invariants hold: (i) $B(u, u.\alpha + 2^j)$ does not satisfy \mathcal{P} ; (ii) $B(u, u.\alpha)$ satisfies \mathcal{P} . The base case is before Step (2). There are two cases: if u is live after round T of Step (5) of the EXPAND, then there is a leader in $B(u, 2^{T+1})$ by the VOTE; otherwise, there is either a fully dormant vertex or a collision in $B(u, 2^{T+1})$. Thus, the invariants hold for the base case.

Now suppose the invariants hold after round j of Step (2). In round $j - 1$, there are two cases. In the first case, $Q(u)$ is set to $Q'(u)$. This means that before Step (2c), $\forall v \in Q(u) = B(u, u.\alpha)$, $B(v, 2^{j-1})$ satisfies \mathcal{P} . Thus, $B(u, u.\alpha + 2^{j-1})$ satisfies \mathcal{P} . Notice that by the induction, $B(u, u.\alpha + 2^j)$ does not satisfy \mathcal{P} . Therefore, the invariants hold after updating $Q(u)$ and $u.\alpha$ in Step (2c). In the second case, $Q(u)$ remains unchanged. There exists $v \in Q(u) = B(u, u.\alpha)$ such that $B(v, 2^{j-1})$ does not satisfy \mathcal{P} which means that $B(u, u.\alpha + 2^{j-1})$ does not satisfy \mathcal{P} . Since $Q(u)$ and $u.\alpha$ can only change together, $B(u, u.\alpha)$ satisfies \mathcal{P} . The invariants also hold. Therefore, after round $j = 0$, $B(u, u.\alpha)$ satisfies \mathcal{P} and $B(u, u.\alpha + 1)$ does not satisfy \mathcal{P} , giving the second part of the lemma. \square

Lemma C.5. For any vertex u , if $u.\beta$ is updated, then $u.\beta = \min_{v:v.l=1} \text{dist}(u, v)$.

Proof. For any vertex u with $u.l = 1$, $u.\beta$ is set to 0. Now, consider a vertex u with $u.l = 0$. By Lemma C.4, $Q(u) = B(u, u.\alpha)$ and there is no leader in $B(u, u.\alpha)$. If $w \in B(u, u.\alpha)$ is marked as a leader-neighbor, there is a vertex v with $v.l = 1$ such that $v \in B(u, u.\alpha + 1)$. Thus, in this case $u.\beta = u.\alpha + 1 = \min_{v:v.l=1} \text{dist}(u, v)$. \square

The following lemma shows that the construction of the tree in Steps (5,6) using the β values is valid.

Lemma C.6. For any vertex u , if $u.\beta > 0$, then there exists an edge $\{u, w\}$ such that $w.\beta = u.\beta - 1$.

Proof. Consider a vertex u with $u.\beta = 1$. Then $u.\alpha = 0$, and by Lemma C.4, $Q(u) = B(u, 0) = \{u\}$. Thus, u is marked as a leader-neighbor which means that there is a graph edge $\{u, v\}$ where v is a leader. Notice that $v.\beta = 0$. So the lemma holds for u with $u.\beta = 1$.

Consider a vertex u with $u.\beta > 1$. By Lemma C.5, there is a leader v such that $\text{dist}(u, v) = u.\beta$. Let $\{u, w\}$ be a graph edge with $\text{dist}(w, v) = u.\beta - 1$, which must exist by Step (5). It suffices to show that $w.\beta = u.\beta - 1$. Since $B(w, u.\beta - 1)$ contains a leader v , and $B(w, u.\beta - 2) \subseteq B(u, u.\beta - 1) = B(u, u.\alpha)$ which does not contain a leader by Lemma C.4, we have that $w.\alpha = u.\beta - 2$. Let $\{x, v\}$ be a graph edge such that $\text{dist}(w, x) = u.\beta - 2$, which must exist by Step (5). Then, x is marked as a leader-neighbor and x is in $B(w, w.\alpha)$. Hence $w.\beta = w.\alpha + 1 = u.\beta - 2 + 1 = u.\beta - 1$. \square

The above lemma implies that if $u.\beta > 0$ then u is a non-root in the next phase due to Steps (5,6):

Corollary C.7. *For any vertex u , if $u.\beta > 0$, then u is finished in the next phase.*

Lemma C.8. *The height of any tree is $O(d)$ after the TREE-LINK.*

Proof. If $u.p$ is updated to w , then $u.\beta = w.\beta + 1$ by Step (6). Thus, the height of a tree is at most $\max_u u.\beta + 1$. By Lemma C.5 and the fact that an ALTER and adding edges never increase the diameter, the height of any tree is at most d . \square

As mentioned at the end of §C.2, by the above lemma and Lemma B.8, the following is immediate:

Corollary C.9. *Each phase of the algorithm takes $O(\log d)$ time.*

Similar to Definition C.1 and Lemma C.3, we can show the following, which guarantees the correctness of Spanning Forest algorithm:

Lemma C.10. *At the end of the TREE-LINK, all the edges \hat{e} with $\hat{e}.f = 1$ constitute a forest. And in each tree of the forest, there is exactly one root.*

C.4 Running Time

Now let us analyze the number of vertices in the next phase. If a vertex u is live after the EXPAND, then it is finished in the next phase. Thus all the vertices in the next phase are dormant after the EXPAND in this phase. Consider a dormant vertex u , and let r be that defined in Definition B.9 with respect to u .

Lemma C.11. *For a dormant non-leader u , if there is a leader in $B(u, r)$, then u is finished in the next phase.*

Proof. Let v be the leader closest to u . Since v is in $B(u, r)$, by the definition of r and Lemma C.4, we have that $u.\alpha = \text{dist}(u, v) - 1$, and $Q(u) = B(u, \text{dist}(u, v) - 1)$ which contains a leader-neighbor. By Step (4b) in the TREE-LINK, $u.\beta = \text{dist}(u, v) > 0$. By Corollary C.7, u is finished in the next phase. \square

Using the above lemma and Corollary C.9, the remaining analysis is almost identical to that in Connected Components algorithm.

Proof of Theorem 2. By Lemma B.11, we have that $|B(u, r)| \leq b^2$ with probability at most b^{-2} . Conditioned on $|B(u, r)| > b^2$, the probability that $B(u, r)$ contains no leader is at most b^{-1} . The number of vertices in the next phase is at most the sum of (i) the number of dormant leaders, (ii) the number of vertices u with $|B(u, r)| \leq b^2$, and (iii) the number of vertices u with $|B(u, r)| > b^2$ and no leader in $B(u, r)$. Thus, the probability that a dormant vertex u is in the next phase is at most $b^{-2/3} + b^{-2} + (1 - b^{-2}) \cdot b^{-1} \leq b^{-1/2}$. By Markov's inequality, the probability of having more than $n^t \cdot b^{-1/4}$ vertices in the next phase is at most $b^{-1/4}$.

Finally, using exactly the same analyses in §B.5 and §B.6, we obtain that the algorithm runs on an ARBITRARY CRCW PRAM and outputs the spanning forest. Moreover, with probability $1 - 1/\text{poly}(m \log n/n)$, the number of phases is $O(\log \log_{m/n} n)$ thus the total running time is $O(\log d \log \log_{m/n} n)$ (cf. Corollary C.9). \square

D Faster Connected Components Algorithm: Full Proof of An $O(\log d + \log \log_{m/n} n)$ Bound

In this section we give the full proof of Theorem 3 by presenting a faster algorithm (see below) for connected components with a detailed analysis.

Faster Connected Components algorithm: COMPACT; repeat {EXPAND-MAXLINK} until the graph has diameter $O(1)$ and all trees are flat; run Connected Components algorithm from §B.

Each iteration of the repeat loop is called a *round*. The *break condition* that the graph has diameter $O(1)$ and all trees are flat shall be tested at the end of each round. The main part of this section is devoted to the repeating EXPAND-MAXLINK. Before that, we explain the method COMPACT.

COMPACT: PREPARE; rename each ongoing vertex to a distinct id in $[2m/\log^c n]$ and assign it a block of size $\max\{m/n, \log^c n\}/\log^2 n$.

Recall that in the PREPARE in §B.2, it is guaranteed that the number of ongoing vertices is at most $m/\log^c n$ with good probability, then we use hashing to assign them blocks. The problem of this method is that in this section, we need, with high probability, all roots to own blocks as we need to union bound (as the events are dependent) all roots on all shortest paths (see §1.2), but the hashing method only gives that each root owns a block with probability $1 - 1/\text{polylog}(n)$. To solve this problem, we need the following tool when implementing COMPACT:

Definition D.1. Given a length- n array A which contains k distinguished elements, approximate compaction is to map all the distinguished elements in A one-to-one to an array of length $2k$.

Lemma D.2 ([Goo91]). There is an ARBITRARY CRCW PRAM algorithm for approximate compaction that runs in $O(\log^* n)$ time and uses $O(n)$ processors with high probability. Moreover, if using $n \log n$ processors, the algorithm runs in $O(1)$ time.¹⁰

Lemma D.3. With good probability, COMPACT renames each ongoing vertex a distinct id in $[2m/\log^c n]$, assigns each of them a block of size $\max\{m/n, \log^c n\}/\log^2 n$, runs in $O(\log \log_{m/n} n)$ time, and uses $O(m)$ processors in total.

Proof. If $m/n > \log^c n$, then the COMPACT does not run Vanilla algorithm and each (ongoing) vertex already has a distinct id in $[n] \subseteq [2m/\log^c n]$. Since there are $\Theta(m)$ processors in total, it is easy to assign each vertex a block of size $m/n/\log^2 n$, which is $\max\{m/n, \log^c n\}/\log^2 n$.

Else if $m/n \leq \log^c n$, the COMPACT runs Vanilla algorithm for $c \log_{8/7} \log n = O(\log \log_{m/n} n)$ phases such that the number of ongoing vertices is at most $m/\log^c n$ with good probability (cf. Corollary B.4). Conditioned on this happening, the COMPACT uses approximate compaction to compact all the at most $m/\log^c n$ ongoing vertices (distinguished elements) to an array of length at most $2m/\log^c n$, whose index in the array serves as a distinct id. By Lemma D.2, this runs in $O(\log^* n)$ time with probability $1 - 1/\text{poly}(n)$. Once the ongoing vertices are indexed, it is easy to assign each of them a block of size $m/(m/\log^c n)/\log^2 n$, which is $\max\{m/n, \log^c n\}/\log^2 n$. By a union bound, the COMPACT takes $O(\log \log_{m/n} n)$ time with good probability. \square

¹⁰[Goo91] actually compacts the k distinguished elements of A to an array of length $(1 + \epsilon)k$ using $O(n/\log^* n)$ processors with probability $1 - 1/c^{n^{1/25}}$ for any constants $\epsilon > 0$ and $c > 1$, but Lemma D.2 suffices for our use. The second part of the lemma is a straightforward corollary (cf. §2.4.1 in [Goo91]).

The renamed vertex id is used only for approximate compaction; for all other cases, we still use the original vertex id. Renaming the vertex id to range $[2m/\log^c n]$ not only facilitates the processor allocation at the beginning, but more importantly, guarantees that each subsequent processor allocation takes $O(1)$ time, because the array A to be compacted has length $|A| \leq 2m/\log^c n$ while the number of processors is at least $|A| \log |A|$ when $c \geq 10$ (cf. Lemma D.2, see §1.2 and details in §D.1).

In the remainder of this section, we present the detailed method EXPAND-MAXLINK in §D.1, then we prove that Faster Connected Components algorithm correctly computes the connected components of the input graph (cf. §D.2), each round (EXPAND-MAXLINK) can be implemented on an ARBITRARY CRCW PRAM in constant time (cf. §D.3), uses $O(m)$ processors over all rounds (cf. §D.4), reduces the diameter to $O(1)$ and flatten all trees in $O(\log d + \log \log_{m/n} n)$ rounds (cf. §D.5), leading to the proof of Theorem 3 in §D.6.

D.1 Algorithmic Framework

In this section, we present the key concepts and algorithmic framework of the EXPAND-MAXLINK.

Level and budget. The *level* $\ell(v)$ of a vertex v is a non-negative integer that can either remain the same or increase by one during a round. By Lemma D.3, at the beginning of round 1, each ongoing vertex v owns a block of size $b_1 := \max\{m/n, \log^c n\} / \log^2 n$, whose level is defined as 1; the level of a non-root vertex whose parent is ongoing is defined as 0 and $b_0 := 0$ for soundness; all other vertices are ignored as their components have been computed. During a round, some roots become non-roots by updating their parents; for those vertex remains to be a root, its level might be increased by one; a root with level ℓ is assigned to a block of size $b_\ell := b_1^{1.01^{\ell-1}}$ at the end of the round. Given $b \geq 0$, a vertex v has *budget* $b(v) := b$ if the maximal-size block owned by v has size b . Each block of size b is partitioned into \sqrt{b} (indexed) tables, each with size \sqrt{b} .

Neighbor set. Recall from §B.2 that the edges that define the current graph include: (i) the (altered) original edges corresponding to edge processors, and (ii) the (altered) added edges in the tables over all rounds of all vertices; any vertex within distance 1 from v (including v) in the current graph is called a neighbor of v . For any vertex v , let $N(v)$ be the set of its neighbors. In Step (3) we use the old $N(v)$ when initializing the loop that enumerates $N(v)$. For any vertex set S , define $N(S) := \bigcup_{w \in S} N(w)$, and define $S.p := \{w.p \mid w \in S\}$.

Hashing. At the beginning of a round, one random hash function h is chosen, then all neighbor roots of all roots use h to do individual hashing in Step (3). As same as in §B, a pairwise independent h suffices, thus each processor only reads two words. The hashing in Step (5) follows the same manner using the same h . For each vertex v , let $H(v)$ be the first table in its block, which will store the added edges incident on v . Step (5) is implemented by storing the old tables for all vertices while hashing new items (copied from the old $H(v)$ and old $H(w)$ in the block of w) into the new table.

EXPAND-MAXLINK:

1. MAXLINK; ALTER.
2. For each root v : increase $\ell(v)$ with probability $10 \log n / b(v)^{0.1}$.
3. For each root v : for each root $w \in N(v)$: if $b(w) = b(v)$ then hash w into $H(v)$.

4. For each root v : if there is a collision in $H(v)$ then mark v as dormant. For each vertex v : if there is a dormant vertex in $H(v)$ then mark v as dormant.
5. For each root v : for each $w \in H(v)$: for each $u \in H(w)$: hash u into $H(v)$. For each root v : if there is a collision in $H(v)$ then mark v as dormant.
6. MAXLINK; SHORTCUT; ALTER.
7. For each root v : if v is dormant and did not increase level in Step (2) then increase $\ell(v)$.
8. For each root v : assign a block of size $b_{\ell(v)}$ to v .

MAXLINK: repeat {for each vertex v : let $u := \arg \max_{w \in N(v).p} \ell(w)$, if $\ell(u) > \ell(v)$ then update $v.p$ to u } for 2 iterations.

The ALTER and SHORTCUT (cf. Steps (1,6)) are the same as in §B. The MAXLINK uses parent links, instead of direct links in §B and §C.

D.2 Correctness

In this section, we prove that Faster Connected Components algorithm correctly computes the connected components of the input graph when it ends.

First of all, we prove a useful property on levels and roots.

Lemma D.4. *If a vertex v is a non-root at any step, then during the execution after that step, v is a non-root, $\ell(v)$ cannot change, and $0 \leq \ell(v) < \ell(v.p)$.*

Proof. The proof is by an induction on rounds. At the beginning of the first round, by the definition of levels (cf. §D.1), each ongoing vertex v has level 1 and $v = v.p$, and each non-root with an ongoing parent has level 0 (vertices in trees rooted at finished vertices are ignored), so the lemma holds.

In Step (1), each iteration of a MAXLINK can only update the parent to a vertex with higher level, which cannot be itself. In Step (2), level increase only applies to roots. The invariant holds after the MAXLINK in Step (6) for the same reasons as above. In SHORTCUT (cf. Step (6)), each vertex v updates its parent to $v.p.p$, which, by the induction hypothesis, must be a vertex with level higher than v if v is a non-root, thus cannot be v . In Step (7), level increase only applies to roots. All other steps and ALTERS do not change the labeled digraph nor levels, giving the lemma. \square

Based on this, we can prove the following key result, which implies the correctness of our algorithm.

Lemma D.5. *The following conditions hold at the end of each round:*

1. *For any component in the input graph, its vertices are partitioned into trees in the labeled digraph such that each tree belongs to exactly one component.*
2. *A tree does not contain all the vertices in its component if and only if there is an edge between a vertex in this tree and a vertex in another tree.*

The proof of Lemma D.5 relies on the following invariant:

Lemma D.6. *For any vertices v and w , if w is a neighbor of v in the current graph or $w = v.p$, then v and w are in the same component. If v is incident with an edge, then there exists a path in the current graph between v and $v.p$.*

Proof. By inspecting Vanilla algorithm (called within COMPACT) and monotonicity, Lemma D.6 holds at the beginning of the first round. By an induction on rounds, it suffices to prove the following: (i) the ALTER (cf. Steps (1,6)) and Step (5) preserve the invariant (v and w are in the same component), as only these two change the neighbor sets; (ii) the MAXLINK and SHORTCUT (cf. Steps (1,6)) preserve the invariant, as only these two change the parents.

Each iteration of the MAXLINK preserves the invariant since both v and its neighbor w , and $w, w.p$ are in the same component by the induction hypothesis. So MAXLINK preserves the invariant. To prove that an ALTER preserves the invariant, let (v', w') be an edge before ALTER such that $v'.p = v$ and $w'.p = w$ (to make w a neighbor of v after an ALTER). By the induction hypothesis, w', w, v', v , and v', w' are in the same component, thus so do v, w . A SHORTCUT preserves the invariant since $v, v.p$ and $v.p, v.p.p$ are both in the same component. Step (5) preserves the invariant, because $H(v), H(w)$ are subsets of the neighbor sets of v and w respectively before Step (5) (cf. Step (3)), thus v, w and w, u are both in the same component by the induction hypothesis, and so do v and its new neighbor u . This finishes the induction and proves the first part of the lemma.

Now we prove the second part of the lemma. By an induction, we assume the invariant holds before some step and prove that each step of EXPAND-MAXLINK preserves the invariant. Before a SHORTCUT, there are paths between $v, v.p$ and $v.p, v.p.p$, thus there is a path between v and $v.p$ after updating $v.p$ to the old $v.p.p$. Before an ALTER, there must exist a child v' of v such that v' is incident with an edge, otherwise v has no incident edge after the ALTER. By the induction hypothesis, there is a path between v' and v , which is altered to be a new path between v and $v.p$ after the ALTER. In each iteration of a MAXLINK, if $v.p$ does not change then the invariant trivially holds; otherwise, let (v, w) be the edge that updates $v.p$ to $w.p$. By the induction hypothesis, there is a path between w and $w.p$, so there is a path between v and $v.p = w.p$ after the iteration. Other steps in the EXPAND-MAXLINK can only add edges, giving the lemma. \square

Proof of Lemma D.5. Conditions (1) and (2) hold at the beginning of the first round by Lemma B.2 and monotonicity of Vanilla algorithm. By Lemma D.4, the levels along any tree path from leaf to root are monotonically increasing, so there cannot be a cycle and the labeled digraph must be a collection of trees. By Lemma D.6 and an induction on tree, all vertices of a tree belong to the same component. This proves Condition (1).

For any tree T , if there is an edge between a vertex v in T and another vertex w in another tree T' , then by Lemma D.6, v, w belong to the same component and so do vertices in T and T' , which means T does not contain all the vertices of this component. On the other hand, assume T does not contain all the vertices in its component, then there must be another tree T' corresponding to the same component. By an induction, we assume the invariant that there is an edge between a vertex v in T and a vertex w in another tree holds, then proves that an ALTER and (each iteration of) a MAXLINK preserve the invariant. An ALTER moves the edge between v and w to their parents in their trees, whose two endpoints must belong to different trees as w not in T , so the invariant holds. Assume a parent update in an iteration of MAXLINK changes the parent of u to a vertex $u.p$ in T' . Let u' be the old parent of u in T . By Lemma D.6, there are paths between u, u' and $u, u.p$, so there must be a path between u' and $u.p$. Since u' and $u.p$ belong to different trees: the new T and T' respectively, there must be an edge from a vertex in T to another tree, and an edge from a vertex in T' to another tree, giving the lemma. \square

Lemma D.7. *If Faster Connected Components algorithm ends, then it correctly computes the connected components of the input graph.*

Proof. When the repeat loop ends, all trees are flat, which means in the last round, all trees are flat before the ALTER in Step (6) since an ALTER does not change the labeled digraph. After the ALTER, all edges are only incident on roots. By Lemma D.5, these trees partition the connected components; moreover, if a root is not the only root in its component then there must be an edge between it and another root in its component,

i.e., the preconditions of Lemma B.2 is satisfied. This implies that the following Connected Components algorithm is applicable and must correctly compute the connected components of the input graph by §B. \square

D.3 Implementation

In this section, we show how to implement EXPAND-MAXLINK such that any of the first $O(\log n)$ rounds runs in constant time with good probability.

Lemma D.8. *With good probability, any of the first $O(\log n)$ rounds can be implemented to run in $O(1)$ time.*

Proof. The ALTER (cf. Steps (1,6)) applies to all edges in the current graph. Since each edge corresponds to a distinct processor, Step (3) and the ALTER take $O(1)$ time.

Steps (2,4,7) and SHORTCUT take $O(1)$ time as each vertex has a corresponding processor and a collision can be detected using the same hash function to write to the same location again.

In each of the 2 iterations in MAXLINK, each vertex v updates its parent to a neighbor parent with the highest level if this level is higher than $\ell(v)$. Since a vertex can increase its level by at most 1 in any round (cf. Steps (2,7)), there are $O(\log n)$ different levels. Let each neighbor of v write its parent with level ℓ to the ℓ -th cell of an array of length $O(\log n)$ in the block of v (arbitrary win). Furthermore, by the definitions of level and budget, the block of any vertex with positive level in any round has size at least $b_1 = \Omega(\log^3 n)$ when $c \geq 10$. (Vertex with level 0 does not have neighbors.) Therefore, we can assign a processor to each pair of the cells in this array such that each non-empty cell can determine whether there is a non-empty cell with larger index then finds the non-empty cell with largest index in $O(1)$ time, which contains a vertex with the maximum level. As a result, Steps (1,6) take $O(1)$ time.

By Step (3), any $w \in H(v)$ has $b(w) = b(v)$, so each $u \in H(w)$ such that $w \in H(v)$ owns a processor in the block of v since $\sqrt{b(v)} \cdot \sqrt{b(w)} = b(v)$ and any vertex in a table is indexed (by its hash value). Therefore, together with collision detection, Step (5) takes $O(1)$ time.

In Step (8), each vertex is assigned to a block. The pool of $\Theta(m)$ processors is partitioned into $\Theta(\log^2 n)$ zones such that the processor allocation in round r for vertices with level ℓ uses the zone indexed by (r, ℓ) , where $r, \ell \in O(\log n)$ in the first $O(\log n)$ rounds. Since there are $\Theta(m)$ processors in total and all the vertex ids are in $[2m/\log^c n]$ with good probability (cf. Lemma D.3), we can use $\Theta(m/\log n)$ processors for each different level and apply Lemma D.2 to index each root in $O(1)$ time with probability $1 - 1/\text{poly}(n)$ such that the indices of vertices with the same level are distinct, then assign each of them a distinct block in the corresponding zone. Therefore, Step (8) takes $O(1)$ time with good probability by a union bound over all $O(\log n)$ levels and rounds.

Finally, we need to implement the break condition in $O(1)$ time, i.e., to determine whether the graph has diameter $O(1)$ and all trees are flat at the end of each round. The algorithm checks the following 2 conditions in each round: (i) all vertices do not change their parents nor labels in this round, and (ii) for any vertices v, w, u such that $w \in H(v)$, $u \in H(w)$ before Step (5), the $h(u)$ -th cell in $H(v)$ already contains u . Conditions (i) and (ii) can be checked in $O(1)$ time by writing a flag to vertex processor v if they do not hold for v , then let each vertex with a flag write the flag to a fixed processor. If there is no such flag then both Conditions (i) and (ii) hold and the loop breaks. If there is a non-flat tree, some parent must change in the SHORTCUT in Step (6). If all trees are flat, they must be flat before the ALTER in Step (6), then an ALTER moves all edges to the roots. Therefore, if Condition (i) holds, all trees are flat and edges are only incident on roots. Moreover, no level changing means no vertex increase its level in Step (2) and there is no dormant vertex in Step (7). So for each root v , $N(v) = H(v)$ after Step (3) and $N(N(v)) = H(v)$ after Step (5) as there is no collision. By Condition (ii), the table $H(v)$ does not change during Step (5), so $N(v) = N(N(v))$. If there exists root v such that there is another root with distance at least 2 from v , then there must exist a vertex $w \neq v$ at distance exactly 2 from v , so $w \notin N(v)$ and $w \in N(N(v))$,

contradicting with $N(v) = N(N(v))$. Therefore, any root is within distance at most 1 from all other roots in its component and the graph has diameter $O(1)$.

Since each of them runs in $O(1)$ time with good probability, the lemma follows. \square

D.4 Number of Processors

In this section, we show that with good probability, the first $O(\log n)$ rounds use $O(m)$ processors in total.

First of all, observe that in the case that a root v with level ℓ increases its level in Step (2) but becomes a non-root at the end of the round, v is not assigned a block of size $b_{\ell(v)}$ in Step (8). Instead, v owns a block of size $b_{\ell} = b_{\ell(v)-1}$ from the previous round. Since in later rounds a non-root never participates in obtaining more neighbors by maintaining its table in Steps (3-5) (which is the only place that requires a larger block), such *flexibility* in the relationship between level and budget is acceptable.

By the fact that any root v at the end of any round owns a block of size $b_{\ell(v)} = b_1^{1.01^{\ell(v)-1}}$, a non-root can no longer change its level nor budget (cf. Lemma D.4), and the discussion above, we obtain:

Corollary D.9. *Any vertex v owns a block of size b at the end of any round where $b_{\ell(v)-1} = b_1^{1.01^{\ell(v)-2}} \leq b \leq b_1^{1.01^{\ell(v)-1}} = b_{\ell(v)}$; if v is a root, then the upper bound on b is tight.*

Secondly, we prove two simple facts related to MAXLINK.

Lemma D.10. *For any vertex v with parent v' and any $w \in N(v)$ before an iteration of MAXLINK, it must be $\ell(w.p) \geq \ell(v')$ after the iteration; furthermore, if $\ell(w.p) > \ell(v)$ before an iteration, then v must be a non-root after the iteration.*

Proof. For any $w \in N(v)$, its parent has level $\max_{u \in N(w)} \ell(u.p) \geq \ell(v')$ after an iteration of MAXLINK. This implies that if $\ell(w.p) > \ell(v)$ before an iteration, then after that $\ell(w.p)$ is at least the level of the old parent of w which is strictly higher than $\ell(v)$, so v must be a non-root. \square

Lemma D.11. *For any root v with budget b at the beginning of any round, if there is a root $w \in N(v)$ with at least $b^{0.1}$ neighbor roots with budget b after Step (1), then v either increases level in Step (2) or is a non-root at the end of the round with probability $1 - n^{-5}$.*

Proof. Assume v does not increase level in Step (2). Let w be any root in $N(v)$ after Step (1). Since each root $u \in N(w)$ with budget b (thus level at least $\ell(v)$ by Corollary D.9) increases its level with probability $10 \log n / b^{0.1}$ independently, with probability at least $1 - (1 - 10 \log n / b^{0.1})^{b^{0.1}} \geq 1 - n^{-5}$, at least one u increases level to at least $\ell(v) + 1$ in Step (2). Since $u \in N(N(v))$ after Step (1), by Lemma D.10, there is a $w' \in N(v)$ such that $\ell(w'.p) \geq \ell(v) + 1$ after the first iteration of MAXLINK in Step (6). Again by Lemma D.10, this implies that v cannot be a root after the second iteration and the following SHORTCUT. \square

Using the above result, we can prove the following key lemma, leading to the total number of processors.

Lemma D.12. *For any root v with budget b at the beginning of any round, $b(v)$ is increased to $b^{1.01}$ in this round with probability at most $n^{-5} + b^{-0.05}$.*

Proof. In Step (2), $\ell(v)$ increases with probability $10 \log n / b^{0.1} \leq b^{-0.08}$ when $c \geq 100$, since $b \geq b_1 \geq \log^{c-2} n$. If $\ell(v)$ does increase here then it cannot increase again in Step (7), so we assume this is not the case (and apply a union bound at the end).

If there is a root $w \in N(v)$ with at least $b^{0.1}$ neighbor roots with budget b after Step (1), then v is a root at the end of the round with probability at most n^{-5} by the assumption and Lemma D.11. So we assume this is not the case.

By the previous assumption we know that at most $b^{0.1}$ vertices are hashed into $H(v)$ in Step (3). By pairwise independency, with probability at most $(b^{0.1})^2/\sqrt{b} = b^{-0.3}$ there is a collision as the table has size \sqrt{b} , which will increase $\ell(v)$ (cf. Steps (4,7)).

Now we assume that there is no collision in $H(v)$ in Step (3), which means $H(v)$ contains all the at most $b^{0.1}$ neighbor roots with budget b . By the same assumption, each such neighbor root w has at most $b^{0.1}$ neighbor roots with budget b , so there is a collision in $H(w)$ in Step (3) with probability at most $(b^{0.1})^2/\sqrt{b} = b^{-0.3}$. By a union bound over all the $|H(v)| \leq b^{0.1}$ such vertices, v is marked as dormant in the second statement of Step (4) (and will increase level in Step (7)) with probability $b^{-0.2}$.

It remains to assume that there is no collision in $H(v)$ nor in any $H(w)$ such that $w \in H(v)$ after Step (4). As each such table contains at most $b^{0.1}$ vertices, in Step (5) there are at most $b^{0.2}$ vertices to be hashed, resulting in a collision in $H(v)$ with probability at most $(b^{0.2})^2/\sqrt{b} = b^{-0.1}$, which increases $\ell(v)$ in Step (7).

Observe that only a root v at the end of the round can increase its budget, and the increased budget must be $b^{1.01}$ since the level can increase by at most 1 during the round and $b = b_{\ell(v)}$ at the beginning of the round by Corollary D.9. By a union bound over the events in each paragraph, $b(v)$ is increased to $b^{1.01}$ with probability at most $b^{-0.08} + n^{-5} + b^{-0.3} + b^{-0.2} + b^{-0.1} \leq n^{-5} + b^{-0.05}$. \square

Finally, we are ready to prove an upper bound on the number of processors.

Lemma D.13. *With good probability, the first $O(\log n)$ rounds use $O(m)$ processors in total.*

Proof. Using Lemma D.12, by a union bound over all $O(n)$ roots, all $O(\log n)$ rounds, and all $O(\log n)$ different budgets (since there are $O(\log n)$ different levels), with probability at least $1 - n^{-3}$, any root v with budget b at the beginning of any round increases its budget to $b^{1.01}$ with probability at most $b^{-0.05}$. We may assume that the $(1 - n^{-3})$ -probability event always holds as a good-probability result follows from a union bound.

The number of processors for (altered) original edges and vertices are clearly $O(m)$ over all rounds (where each vertex processor needs $O(1)$ private memory to store the corresponding parent, (renamed) vertex id, hash function, level, and budget). By the proof of Lemma D.8, with good probability the processor allocation in Step (8) always succeeds and has a multiplicative factor of 2 in the number of processors allocated before mapping the indexed vertices to blocks (cf. Definition D.1). Therefore, we only need to bound the number of processors in blocks that are assigned to a vertex in Step (8) in all $O(\log n)$ rounds.

For any positive integer ℓ , let n_ℓ be the number of vertices that ever reaches budget b_ℓ during the first $O(\log n)$ rounds. For any vertex v that ever reaches budget b_ℓ , it has exactly one chance to reach budget $b_{\ell+1}$ in a round if v is a root in that round, which happens with probability at most $b_\ell^{-0.05}$. By a union bound over all $O(\log n)$ rounds, v reaches budget $b_{\ell+1}$ with probability at most $O(\log n) \cdot b_\ell^{-0.05} \leq b_\ell^{-0.04}$ when $c \geq 200$, since $b_\ell \geq b_1 \geq \log^{c-2} n$. We obtain $\mathbb{E}[n_{\ell+1} | n_\ell] \leq n_\ell \cdot b_\ell^{-0.04}$, thus by $b_{\ell+1} = b_\ell^{1.01}$, it must be:

$$\mathbb{E}[n_{\ell+1} b_{\ell+1} | n_\ell] \leq n_\ell \cdot b_\ell^{-0.04} \cdot b_\ell^{1.01} = n_\ell b_\ell \cdot b_\ell^{-0.03}.$$

By Markov's inequality, $n_{\ell+1} b_{\ell+1} \leq n_\ell b_\ell$ with probability at least $1 - b_\ell^{-0.03} \geq 1 - b_1^{-0.03}$. By a union bound on all $\ell \in O(\log n)$, $n_\ell b_\ell \leq n_1 b_1$ for all $\ell \in O(\log n)$ with probability at least $1 - O(\log n) \cdot b_1^{-0.03} \geq 1 - b_1^{-0.01}$, which is $1 - 1/\text{poly}((m \log n)/n)$ by $b_1 = \max\{m/n, \log^c n\}/\log^2 n$ and $c \geq 100$. So the number of new allocated processors for vertices with any budget in any of the first $O(\log n)$ rounds is at most $n_1 b_1$ with good probability.

Recall from Lemma D.3 that with good probability,

$$\begin{aligned} n_1 \cdot b_1 &= \min\{n, 2m/\log^c n\} \cdot \max\{m/n, \log^c n\}/\log^2 n \\ &= O(m/\log^2 n). \end{aligned}$$

Finally, by a union bound over all the $O(\log n)$ levels and $O(\log n)$ rounds, with good probability the total number of processors is $O(m)$. \square

D.5 Diameter Reduction

Let $R := O(\log d + \log \log_{m/n} n)$ where the constant hidden in O will be determined later in this section. The goal is to prove that $O(R)$ rounds of EXPAND-MAXLINK suffice to reduce the diameter of the graph to $O(1)$ and flatten all trees with good probability.

In a high level, our algorithm/proof is divided into the following 3 stages/lemmas:

Lemma D.14. *With good probability, after round R , the diameter of the graph is $O(R)$.*

Lemma D.15. *With good probability, after round $O(R)$, the diameter of the graph is at most 1.*

Lemma D.16. *With good probability, after round $O(R)$, the diameter of the graph is $O(1)$ and all trees are flat.*

Lemma D.16 implies that the repeat loop must end in $O(R)$ rounds with good probability by the break condition.

D.5.1 Path Construction

To formalize and quantify the effect of reducing the diameter, consider any shortest path P in the original input graph G . When (possibly) running Vanilla algorithm on G in COMPACT, an ALTER replaces each vertex on P by its parent, resulting in a path P' of the same length as P . (Note that P' might not be a shortest path in the current graph and can contain loops.) Each subsequent ALTER in COMPACT continues to replace the vertices on the old path by their parents to get a new path. Define P_1 to be the path obtained from P by the above process at the beginning of round 1, by $|P| \leq d$ we immediately get:

Corollary D.17. *For any P_1 corresponding to a shortest path in G , $|P_1| \leq d$.*

In the following, we shall fix a shortest path P in the original graph and consider its corresponding paths over rounds. Each ALTER (cf. Steps (1,6)) in each round does such replacements to the old paths, but we also add edges to the graph for reducing the diameter of the current graph: for any vertices v and w on path P' , if the current graph contains edge (v, w) , then all vertices exclusively between v and w can be removed from P' , which still results in a valid path in the current graph from the first to the last vertex of P' , reducing the length. If all such paths reduce their lengths to at most d' , the diameter of the current graph is at most d' . Formally, we have the following inductive construction of paths for diameter reduction:

Definition D.18 (path construction). *Let all vertices on P_1 be active. For any positive integer r , given path P_r with at least 4 active vertices at the beginning of round r , EXPAND-MAXLINK constructs P_{r+1} by the following 7 phases:*

1. *The ALTER in Step (1) replaces each vertex v on P_r by $v' := v.p$ to get path $P_{r,1}$. For any v' on $P_{r,1}$, let \underline{v}' be on P_r such that $\underline{v}'.p = v'$.¹¹*
2. *Let the subpath containing all active vertices on $P_{r,1}$ be $P_{r,2}$.*
3. *After Step (5), set i as 1, and repeat the following until $i \geq |P_{r,2}| - 1$: let $v' := P_{r,2}(i)$, if \underline{v}' is a root and does not increase level during round r then: if the current graph contains edge $(v', P_{r,2}(i+2))$ then mark $P_{r,2}(i+1)$ as skipped and set i as $i+2$; else set i as $i+1$.*
4. *For each $j \in [i+1, |P_{r,2}|+1]$, mark $P_{r,2}(j)$ as passive.*

¹¹Semantically, there might be more than one \underline{v}' with the same parent v' , but our reference always comes with an index on the replaced path, whose corresponding position in the old path is unique. All subsequent names follow this manner.

5. Remove all skipped and passive vertices from $P_{r,2}$ to get path $P_{r,5}$.
6. Concatenate $P_{r,5}$ with all passive vertices on $P_{r,1}$ and $P_{r,2}$ to get path $P_{r,6}$.
7. The ALTER in Step (6) replaces each vertex v on $P_{r,6}$ by $v.p$ to get path P_{r+1} .

For any vertex v on P_r that is replaced by v' in Phase (1), if v' is not skipped in Phase (3), then let \bar{v} be the vertex replacing v' in Phase (7), and call \bar{v} the corresponding vertex of v in round $r + 1$.

Lemma D.19. *For any non-negative integer r , the P_{r+1} constructed in Definition D.18 is a valid path in the graph and all passive vertices are consecutive from the successor of the last active vertex to the end of P_{r+1} .*

Proof. The proof is by an induction on r . Initially, P_1 is a valid path by our discussion on ALTER at the beginning of this section: it only replaces edges by new edges in the altered graph; moreover, the second part of the lemma is trivially true as all vertices are active. Assuming P_r is a valid path and all passive vertices are consecutive from the successor of the last active vertex to the end of the path. We show the inductive step by proving the invariant after each of the 7 phases in Definition D.18. Phase (1) maintains the invariant. In Phase (2), $P_{r,2}$ is a valid path as all active vertices are consecutive at the beginning of $P_{r,1}$ (induction hypothesis). In Phase (3), if a vertex v is skipped, then there is an edge between its predecessor and successor on the path; otherwise there is an edge between v and its successor by the induction hypothesis; all passive vertices are consecutive from the successor of the last non-skipped vertex to the end of $P_{r,2}$ (cf. Phase (4)), so the invariant holds. In Phase (6), since the first passive vertex on $P_{r,2}$ is a successor of the last vertex on $P_{r,5}$ and the last passive vertex on $P_{r,2}$ is a predecessor of the first passive vertex on $P_{r,1}$ (induction hypothesis), the invariant holds. Phase (7) maintains the invariant. Therefore, P_{r+1} is a valid path and all passive vertices are consecutive from the successor of the last active vertex to the end of P_{r+1} . \square

Now we relate the path construction to the diameter of the graph:

Lemma D.20. *For any positive integer r , the diameter of the graph at the end of round r is $O(\max_{P_r} |P_{r,2}| + r)$.*

Proof. Any shortest path P in the original input graph from s to t is transformed to the corresponding P_1 at the beginning of round 1, which maintains connectivity between the corresponding vertices of s and t on P_1 respectively. By an induction on the number of ALTERS and Lemma D.19, the corresponding vertices of s and t are still connected by path P_{r+1} at the end of round r . Note that by Lemma D.19, P_{r+1} can be partitioned into two parts after Phase (2): subpath $P_{r,2}$ and the subpath containing only passive vertices. Since in each round we mark at most 2 new passive vertices (cf. Phases (3,4)), we get $|P_{r+1}| \leq |P_{r,5}| + 2r \leq |P_{r,2}| + 2r$. If any path P_{r+1} that corresponds to a shortest path in the original graph have length at most d' , the graph at the end of round r must have diameter at most d' , so the lemma follows. \square

It remains to bound the length of any $P_{r,2}$ in any round r , which relies on the following potential function:

Definition D.21. *For any vertex v on P_1 , define its potential $\phi_1(v) := 1$. For any positive integer r , given path P_r with at least 4 active vertices at the beginning of round r and the potentials of vertices on P_r , define the potential of each vertex on P_{r+1} based on Definition D.18 as follows: ¹²*

- For each v replaced by $v.p$ in Phase (1), $\phi_{r,1}(v.p) := \phi_r(v)$.

¹²The second subscript of a potential indicates the phase to obtain that potential; the subscript for paths in Definition D.18 follows the same manner.

- After Phase (4), for each active vertex v on $P_{r,2}$, if the successor w of v is skipped or passive, then $\phi_{r,4}(v) := \phi_{r,1}(v) + \phi_{r,1}(w)$.
- After Phase (6), for each vertex v on $P_{r,6}$, if v is active on $P_{r,2}$, then $\phi_{r,6}(v) := \phi_{r,4}(v)$, otherwise $\phi_{r,6}(v) := \phi_{r,1}(v)$.
- For each v replaced by $v.p$ in Phase (7), $\phi_{r+1}(v.p) := \phi_{r,6}(v)$.

We conclude this section by some useful properties of potentials.

Lemma D.22. *For any path P_r at the beginning of round $r \geq 1$, the following holds: (i) $\sum_{v \in P_r} \phi_r(v) \leq d + 1$; (ii) for any v on P_r , $\phi_r(v) \geq 1$; (iii) for any non-skipped v on $P_{r,2}$ and its corresponding vertex \bar{v} on P_{r+1} , $\phi_{r+1}(\bar{v}) \geq \phi_r(v)$.*

Proof. The proof is by an induction on r . The base case follows from $\phi(v) = 1$ for each v on P_1 (cf. Definition D.21) and Corollary D.17. For the inductive step, note that by Definition D.21, the potential of a corresponding vertex is at least the potential of the corresponding vertex in the previous round (and can be larger in the case that its successor is skipped or passive). This gives (ii) and (iii) of the lemma. For any vertex u on $P_{r,2}$, if both u and its successor are active, then $\phi_r(u)$ is presented for exactly 1 time in $\sum_{v \in P_r} \phi_r(v)$ and $\sum_{v \in P_{r+1}} \phi_{r+1}(v)$ respectively; if u is active but its successor w is skipped or passive, then $\phi_r(u) + \phi_r(w)$ is presented for exactly 1 time in each summations as well; if u and its predecessor are both passive, then $\phi_r(u)$ is presented only in $\sum_{v \in P_r} \phi_r(v)$; the potential of the last vertex on $P_{r,2}$ might not be presented in $\sum_{v \in P_{r+1}} \phi_{r+1}(v)$ depending on i after Phase (3). Therefore, $\sum_{v \in P_{r+1}} \phi_{r+1}(v) \leq \sum_{v \in P_r} \phi_r(v)$ and the lemma holds. \square

D.5.2 Remaining Proofs: Proof of Lemma D.14

Now we prove Lemma D.14, which relies on several results. First of all, we need an upper bound on the maximal possible level:

Lemma D.23. *With good probability, the level of any vertex in any of the first $O(\log n)$ rounds is at most $L := 1000 \max\{2, \log_{m/n} n\}$.*

Proof. By Lemma D.13, with good probability the total number of processors used in the first $O(\log n)$ rounds is $O(m)$. We shall condition on this happening then assume for contradiction that there is a vertex v with level at least L in some round.

If $\log \log_{m/n} n \leq 2$, then $m/n \geq n^{1/4}$. By Corollary D.9, a block owned by v has size at least

$$b_1^{1.01^{2000-2}} \geq b_1^{20} \geq (m/n / \log^2 n)^{20} \geq (n^{1/5})^{20} = n^4 \geq m^2,$$

which is a contradiction as the size of this block owned by v exceeds the total number of processors $O(m)$.

Else if $\log \log_{m/n} n > 2$, then by Corollary D.9, a block owned by v has size at least

$$b_1^{1.01^{L-2}} \geq b_1^{1.01^{999 \log \log_{m/n} n}} \geq b_1^{(\log_{m/n} n)^{10}} \geq b_1^{8 \log_{m/n} n}. \quad (4)$$

Whether $m/n \leq \log^c n$ or not, if $c \geq 10$, it must be $b_1 = \max\{m/n, \log^c n\} / \log^2 n \geq \sqrt{m/n}$. So the value of (4) is at least $n^4 \geq m^2$, contradiction. Therefore, the level of any vertex is at most L . \square

We also require the following key lemma:

Lemma D.24. *For any root v and any $u \in N(N(v))$ at the beginning of any round, let u' be the parent of u after Step (1). If v does not increase level and is a root during this round, then $u' \in H(v)$ after Step (5).*

To prove Lemma D.24, we use another crucial property of the algorithm, which is exactly the reason behind the design of MAXLINK.

Lemma D.25. *For any root v and any $u \in N(N(v))$ at the beginning of any round, if v does not increase level in Step (2) and is a root at the end of the round, then $u.p$ is a root with budget $b(v)$ after Step (1).*

Proof. By Lemma D.4, v is a root during this round. For any $w \in N(v)$ and any $u \in N(w)$, applying Lemma D.10 for 2 times, we get that $\ell(v) \leq \ell(w.p)$ and $\ell(v) \leq \ell(u.p)$ after the MAXLINK in Step (1). If there is a $u \in N(N(v))$ such that $u.p$ is a non-root or $\ell(u.p) > \ell(v)$ before the ALTER in Step (1), it must be $\ell(u.p.p) > \ell(v)$ by Lemma D.4. Note that $u.p$ is in $N(N(v))$ after the ALTER, which still holds before Step (6) as we only add edges. By Lemma D.10, there is a $w' \in N(v)$ such that $\ell(w'.p) > \ell(v)$ after the first iteration of MAXLINK in Step (6). Again by Lemma D.10, this implies that v cannot be a root after the second iteration, a contradiction. Therefore, for any $u \in N(N(v))$, $u.p$ is a root with level $\ell(v)$ (thus budget $b(v)$) after Step (1). \square

With the help of Lemma D.25 we can prove Lemma D.24:

Proof of Lemma D.24. For any vertex u , let $N'(u)$ be the set of neighbors after Step (1). First of all, we show that after Step (5), $H(v)$ contains all roots in $N'(N'(v))$ with budget b , where b is the budget of v at the beginning of the round. For any root $w \in N'(v)$, in Step (3), all roots with budget $b(w)$ in $N'(w)$ are hashed into $H(w)$. If there is a collision in any $H(w)$, then v must be dormant (cf. Step (4)) thus increases level in either Step (2) or (7), contradiction. So there is no collision in $H(w)$ for any $w \in N'(v)$, which means $H(w) \supseteq N'(w)$. Recall that $v \in N'(v)$ and we get that all roots with budget $b(w) = b$ from $N'(N'(v))$ are hashed into $H(v)$ in Step (5). Again, if there is a collision, then v must be dormant and increase level in this round. Therefore, $N'(N'(v)) \subseteq H(v)$ at the end of Step (5).

By Lemma D.25, for any $u \in N(N(v))$ at the beginning of any round, $u' = u.p$ is a root with budget b in $N'(N'(v))$ after Step (1). Therefore, $u' \in H(v)$ at the end of Step (5), giving Lemma D.24. \square

The proof of Lemma D.14 relies on the following lemma based on potentials:

Lemma D.26. *At the beginning of any round $r \geq 1$, for any active vertex v on any path P_r , $\phi_r(v) \geq 2^{r-\ell(v)}$.*

Proof. The proof is by an induction on r . The base case holds because for any (active) vertex v on P_1 , $\phi_1(r) = 1$ and $r = \ell(v) = 1$ (other vertices do not have incident edges). Now we prove the inductive step from r to $r + 1$ given that the corresponding vertex \bar{v} of $v \in P_r$ is on P_{r+1} and active.

Suppose v is a non-root at the end of round r . If v is a non-root at the end of Step (1), then $\ell(v.p) > \ell(v)$ after Step (1) by Lemma D.4, and $\ell(\bar{v}) \geq \ell(v.p) > \ell(v)$; else if v first becomes a non-root in Step (6), then $\bar{v} = v.p$ and $\ell(v.p) > \ell(v)$ after Step (6) by Lemma D.4. So by the induction hypothesis, $\phi_{r+1}(\bar{v}) \geq \phi_r(v) \geq 2^{r-\ell(v)} \geq 2^{r+1-\ell(\bar{v})}$.

Suppose v increases its level in round r . Let ℓ be the level of v at the beginning of round r . If the increase happens in Step (2), then v is a root after Step (1). Whether v changes its parent in Step (6) or not, the level of $\bar{v} = v.p$ is at least $\ell + 1$. Else if the increase happens in Step (7), then v is a root after Step (6). So $\bar{v} = v$ and its level is at least $\ell + 1$ at the end of the round. By the induction hypothesis, $\phi_{r+1}(\bar{v}) \geq \phi_r(v) \geq 2^{r-\ell} \geq 2^{r+1-\ell(\bar{v})}$.

It remains to assume that v is a root and does not increase level during round r . By Lemma D.24, for any $u \in N(N(v))$ at the beginning of round r , the parent u' of u after Step (1) is in $H(v)$ after Step (5). Since v is a root during the round, it remains on $P_{r,2}$ after Phase (2). We discuss two cases depending on whether v is at position before $|P_{r,2}| - 1$ or not.

In Phase (3), note that if $v = P_{r,2}(i)$ where $i < |P_{r,2}| - 1$, then $P_{r,2}(i + 2)$ is the parent of a vertex in $N(N(v))$ after Step (1), which must be in $H(v)$ after Step (5). Therefore, the graph contains edge

$(v, P_{r,2}(i+2))$ and $v' := P_{r,2}(i+1)$ is skipped, thus $\phi_{r,4}(v) = \phi_{r,1}(v) + \phi_{r,1}(v')$ by Definition D.21. Since $i+1 \leq |P_{r,2}| + 1$, v' is an active vertex on $P_{r,1}$. By the induction hypothesis, $\phi_{r,1}(v') = \phi_r(\underline{v}') \geq 2^{r-\ell(\underline{v}')$ (recall that \underline{v}' is replaced by its parent v' in Phase (1)/Step (1)). If $\ell(\underline{v}') > \ell(v)$, then applying Lemma D.10 for 2 times we get that v is a non-root after Step (1), a contraction. Therefore, $\phi_{r,1}(v') \geq 2^{r-\ell(\underline{v}')} \geq 2^{r-\ell(v)}$ and $\phi_{r,4}(v) \geq \phi_{r,1}(v) + \phi_{r,1}(v') \geq \phi_r(v) + 2^{r-\ell(v)} \geq 2^{r+1-\ell(v)}$.

On the other hand, if $i \geq |P_{r,2}| - 1$ is reached after Phase (3), it must be $i < |P_{r,2}| + 1$ by the break condition of the loop in Phase (3). Note that $v' := P_{r,2}(i+1)$ is marked as passive in Phase (4), and by Definition D.21, $\phi_{r,4} = \phi_{r,1}(v) + \phi_{r,1}(v')$. Moreover, since $i+1 \leq |P_{r,2}| + 1$, v' is an active vertex on $P_{r,1}$. Using the same argument in the previous paragraph, we obtain $\phi_{r,4}(v) \geq 2^{r+1-\ell(v)}$.

By Definition D.21, after Phase (7), $\phi_{r+1}(\bar{v}) = \phi_{r,6}(v) = \phi_{r,4}(v) \geq 2^{r+1-\ell(v)} = 2^{r+1-\ell(\bar{v})}$. As a result, the lemma holds for any active vertex \bar{v} on P_{r+1} , finishing the induction and giving the lemma. \square

Proof of Lemma D.14. Let $R := \log d + L$, where L is defined in Lemma D.23. By Lemma D.23, with good probability, $\ell(v) \leq L$ for any vertex v in any of the first $O(\log n)$ rounds, and we shall condition on this happening. By Lemma D.26, at the beginning of round R , if there is a path P_R of at least 4 active vertices, then for any of these vertices v , it must be $\phi_R(v) \geq 2^{R-\ell(v)} \geq 2^{R-L} \geq d$. So $\sum_{v \in P_R} \phi_r(v) \geq 4d > d + 1$, contradicting with Lemma D.22. Thus, any path P_R has at most 3 active vertices, which means $|P_{R,2}| \leq 3$ by Definition D.18. Therefore, by Lemma D.20, the diameter of the graph at the end of round R is $O(R)$ with good probability. \square

D.5.3 Remaining Proofs: Proof of Lemma D.16

Now we prove Lemma D.16 as outlined at the beginning of §D.5. Based on the graph and any P_R at the beginning of round $R + 1$, we need a (much simpler) path construction:

Definition D.27. For any integer $r > R$, given path P_r with $|P_r| \geq 3$ at the beginning of round r , EXPAND-MAXLINK constructs P_{r+1} by the following:

1. The ALTER in Step (1) replaces each vertex v on P_r by $v' := v.p$ to get path $P_{r,1}$. For any v' on $P_{r,1}$, let \underline{v}' be on P_r such that $\underline{v}.p = v'$.
2. After Step (5), let $v' := P_{r,1}(1)$, if \underline{v}' is a root at the end of round r and does not increase level during round r then: if the current graph contains edge $(v', P_{r,1}(3))$ then remove $P_{r,1}(2)$ to get path $P_{r,2}$.
3. The ALTER in Step (6) replaces each vertex v on $P_{r,2}$ by $v.p$ to get path P_{r+1} .

For any vertex v on P_r that is replaced by v' in the first step, if v' is not removed in the second step, then let \bar{v} be the vertex replacing v' in the third step, and call \bar{v} the corresponding vertex of v in round $r + 1$.

An analog of Lemma D.19 immediately shows that P_r is a valid path for any $r \geq R + 1$. The proof of Lemma D.15 is simple enough without potential:

Proof of Lemma D.15. By Lemma D.14, at the beginning of round $R + 1$, with good probability, any P_{R+1} has length $O(R)$. We shall condition on this happening and apply a union bound at the end of the proof. In any round $r > R$, for any path P_r with $|P_r| \geq 3$, consider the first vertex v' on $P_{r,1}$ (cf. Definition D.27). If \underline{v}' is a non-root or increases its level during round r , then by the first 3 paragraphs in the proof of Lemma D.26, it must be $\ell(\bar{v}') \geq \ell(\underline{v}') + 1$. Otherwise, by Lemma D.24, there is an edge between v' and the successor of its successor in the graph after Step (5), which means the successor of v' on $P_{r,1}$ is removed in the second step of Definition D.27. Therefore, the number of vertices on P_{r+1} is one less than P_r if $\ell(\bar{v}') = \ell(\underline{v}')$ as the level of a corresponding vertex cannot be lower. By Lemma D.23, with good probability, the level of any vertex in any of the $O(\log n)$ rounds cannot be higher than L . As P_{R+1} has

$O(R)$ vertices, in round $r = O(R) + L + R = O(R) \leq O(\log n)$, the number of vertices on any P_r is at most 2. Therefore, the diameter of the graph after $O(R)$ rounds is at most 1 with good probability. \square

Proof of Lemma D.16. Now we show that after the diameter reaches 1, if the loop has not ended, then the loop must break in $2L + \log_{5/4} L$ rounds with good probability, i.e., the graph has diameter $O(1)$ and all trees are flat.

For any component, let u be a vertex in it with the maximal level and consider any (labeled) tree of this component. For any vertex v in this tree that is incident with an edge, since the diameter is at most 1, v must have an edge with u , which must be a root. So v updates its parent to a root with the maximal level after a MAXLINK, then any root must have the maximal level in its component since a root with a non-maximal level before the MAXLINK must have an edge to another tree (cf. Lemma D.5). Moreover, if v is a root, this can increase the maximal height among all trees in its component by 1.

Consider the tree with maximal height ξ in the labeled digraph after Step (1). By Lemmas D.4 and D.23, with good probability $\xi \leq L$. The maximal level can increase by at most 1 in this round. If it is increased in Step (7), the maximal height is at most $\lceil \xi/2 \rceil + 1$ after the MAXLINK in the next round; otherwise, the maximal height is at most $\lceil (\xi + 1)/2 \rceil \leq \lceil \xi/2 \rceil + 1$. If $\xi \geq 4$, then the maximal height of any tree after Step (1) in the next round is at most $(4/5)\xi$ (the worst case is that a tree with height 5 gets shortcut to height 3 in Step (6) and increases its height by 1 in the MAXLINK of Step (1) in the next round). Therefore, after $\log_{5/4} L$ rounds, the maximal height of any tree is at most 3.

Beyond this point, if any tree has height 1 after Step (1), then it must have height 1 at the end of the previous round since there is no incident edge on leaves after the ALTER in the previous round, thus the loop must have been ended by the break condition. Therefore, the maximal-height tree (with height 3 or 2) cannot increase its height beyond this point. Suppose there is a tree with height 3, then if the maximal level of vertices in this component does not change during the round, this tree cannot increase its height in the MAXLINK of Step (6) nor that of Step (1) in the next round, which means it has height at most 2 as we do a SHORTCUT in Step (6). So after L rounds, all trees have heights at most 2 after Step (1). After that, similarly, if the maximal level does not increase, all trees must be flat. Therefore, after additional L rounds, all trees are flat after Step (1). By the same argument, the loop must have ended in the previous round. The lemmas follows immediately from $L = O(R)$ and Lemma D.15. \square

D.6 Proof of Theorem 3

With all pieces from §D.2-§D.5 and Theorem 1, we are ready to prove Theorem 3.

Proof of Theorem 3. By Lemma D.16 and the break condition, the repeat loop in Faster Connected Components algorithm runs for $O(\log d + \log \log_{m/n} n) \leq O(\log n)$ rounds. So by Lemma D.8 and a union bound, this loop runs in $O(\log d + \log \log_{m/n} n)$ time with good probability. Moreover, since the diameter is $O(1)$ after the loop (cf. Lemma D.16), the following Connected Component algorithm runs in $O(\log \log_{m/n} n)$ time with good probability (cf. Theorem 1). By Lemma D.3, the method COMPACT also runs in $O(\log \log_{m/n} n)$ time with good probability. Therefore, by a union bound, the total running time is $O(\log d + \log \log_{m/n} n)$ with good probability. When it ends, by Lemma D.7, the algorithm correctly computes the connected components of the input graph. Since there are at most $O(\log n)$ rounds in the loop, by Lemma D.3, Lemma D.13, and Theorem 1, the total number of processors is $O(m)$ with good probability by a union bound. Theorem 3 follows from a union bound at last. \square

References

- [ABSS15] Michael Axtmann, Timo Bingmann, Peter Sanders, and Christian Schulz. Practical massively parallel sorting. In *Proceedings of the 27th ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2015, Portland, OR, USA, June 13-15, 2015*, pages 13–23, 2015.
- [Ajt90] Miklós Ajtai. Approximate counting with uniform constant-depth circuits. In *Advances In Computational Complexity Theory, Proceedings of a DIMACS Workshop, New Jersey, USA, December 3-7, 1990*, pages 1–20, 1990.
- [Akl89] Selim G. Akl. *Design and analysis of parallel algorithms*. Prentice Hall, 1989.
- [AS87] Baruch Awerbuch and Yossi Shiloach. New connectivity and MSF algorithms for shuffle-exchange network and PRAM. *IEEE Trans. Computers*, 36(10):1258–1263, 1987.
- [AS17] Michael Axtmann and Peter Sanders. Robust massively parallel sorting. In *Proceedings of the Nineteenth Workshop on Algorithm Engineering and Experiments, ALENEX 2017, Barcelona, Spain, Hotel Porta Fira, January 17-18, 2017*, pages 83–97, 2017.
- [ASS⁺18] Alexandr Andoni, Zhao Song, Clifford Stein, Zhengyu Wang, and Peilin Zhong. 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*, pages 674–685, 2018.
- [ASW19] Sepehr Assadi, Xiaorui Sun, and Omri Weinstein. Massively parallel algorithms for finding well-connected components in sparse graphs. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 461–470, 2019.
- [BDE⁺19] Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Lacki, and Vahab S. Mirrokni. Near-optimal massively parallel graph connectivity. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 1615–1636, 2019.
- [BH89] Paul Beame and Johan Håstad. Optimal bounds for decision problems on the CRCW PRAM. *J. ACM*, 36(3):643–670, 1989.
- [BKS17] Paul Beame, Paraschos Koutris, and Dan Suciu. Communication steps for parallel query processing. *J. ACM*, 64(6):40:1–40:58, 2017.
- [CDR86] Stephen A. Cook, Cynthia Dwork, and Rüdiger Reischuk. Upper and lower time bounds for parallel random access machines without simultaneous writes. *SIAM J. Comput.*, 15(1):87–97, 1986.
- [DG08] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.
- [DKR94] Martin Dietzfelbinger, Mirosław Kutyłowski, and Rüdiger Reischuk. Exact lower time bounds for computing boolean functions on CREW prams. *J. Comput. Syst. Sci.*, 48(2):231–254, 1994.
- [FMRW85] Faith E. Fich, Friedhelm Meyer auf der Heide, Prabhakar Ragde, and Avi Wigderson. One, two, three \dots infinity: Lower bounds for parallel computation. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, pages 48–58, 1985.

- [Gaz91] Hillel Gazit. An optimal randomized parallel algorithm for finding connected components in a graph. *SIAM J. Comput.*, 20(6):1046–1067, 1991.
- [GKP94] Steve Goddard, Subodh Kumar, and Jan F. Prins. Connected components algorithms for mesh-connected parallel computers. In *Parallel Algorithms, Proceedings of a DIMACS Workshop, Brunswick, New Jersey, USA, October 17-18, 1994*, pages 43–58, 1994.
- [GMV91] Joseph Gil, Yossi Matias, and Uzi Vishkin. Towards a theory of nearly constant time parallel algorithms. In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*, pages 698–710, 1991.
- [GMW90] Joseph Gil, Friedhelm Meyer auf der Heide, and Avi Wigderson. Not all keys can be hashed in constant time (preliminary version). In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 244–253, 1990.
- [Goo91] Michael T. Goodrich. Using approximation algorithms to design parallel algorithms that may ignore processor allocation (preliminary version). In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*, pages 711–722, 1991.
- [Gre94] John Greiner. A comparison of parallel algorithms for connected components. In *Proceedings of the sixth annual ACM symposium on Parallel algorithms and architectures, SPAA 1994*, pages 16–25. ACM, 1994.
- [GSZ11] Michael T. Goodrich, Nodari Sitchinava, and Qin Zhang. Sorting, searching, and simulation in the mapreduce framework. In *Algorithms and Computation - 22nd International Symposium, ISAAC 2011, Yokohama, Japan, December 5-8, 2011. Proceedings*, pages 374–383, 2011.
- [HCS79] Daniel S. Hirschberg, Ashok K. Chandra, and Dilip V. Sarwate. Computing connected components on parallel computers. *Commun. ACM*, 22(8):461–464, 1979.
- [HRD97] Tsan-Sheng Hsu, Vijaya Ramachandran, and Nathaniel Dean. Parallel implementation of algorithms for finding connected components in graphs. *Parallel Algorithms: Third DIMACS Implementation Challenge, October 17-19, 1994*, 30:20, 1997.
- [HZ96] Shay Halperin and Uri Zwick. An optimal randomised logarithmic time connectivity algorithm for the EREW PRAM. *J. Comput. Syst. Sci.*, 53(3):395–416, 1996.
- [HZ01] Shay Halperin and Uri Zwick. Optimal randomized erew pram algorithms for finding spanning forests. *Journal of Algorithms*, 39(1):1–46, 2001.
- [JM97] Donald B. Johnson and Panagiotis Takis Metaxas. Connected components in $O(\log^{3/2} n)$ parallel time for the CREW PRAM. *J. Comput. Syst. Sci.*, 54(2):227–242, 1997.
- [KSV10] Howard J. Karloff, Siddharth Suri, and Sergei Vassilvitskii. 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*, pages 938–948, 2010.
- [LT19] S. Cliff Liu and Robert E. Tarjan. Simple concurrent labeling algorithms for connected components. In *2nd Symposium on Simplicity in Algorithms, SOSA@SODA 2019, January 8-9, 2019 - San Diego, CA, USA*, pages 3:1–3:20, 2019.
- [MR85] Gary L. Miller and John H. Reif. Parallel tree contraction and its application. In *26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, USA, 21-23 October 1985*, pages 478–489, 1985.

- [Rei84] John H Reif. Optimal parallel algorithms for graph connectivity. Technical report, HARVARD UNIV CAMBRIDGE MA AIKEN COMPUTATION LAB, 1984.
- [RVW18] Tim Roughgarden, Sergei Vassilvitskii, and Joshua R. Wang. Shuffles and circuits (on lower bounds for modern parallel computation). *J. ACM*, 65(6):41:1–41:24, 2018.
- [SDB14] Julian Shun, Laxman Dhulipala, and Guy E. Blelloch. A simple and practical linear-work parallel algorithm for connectivity. In *26th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '14, Prague, Czech Republic - June 23 - 25, 2014*, pages 143–153, 2014.
- [SRT18] Stergios Stergiou, Dipen Rughwani, and Kostas Tsioutsoulouklis. Shortcutting label propagation for distributed connected components. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM 2018, Marina Del Rey, CA, USA, February 5-9, 2018*, pages 540–546, 2018.
- [SV82] Yossi Shiloach and Uzi Vishkin. An $O(\log n)$ parallel connectivity algorithm. *J. Algorithms*, 3(1):57–67, 1982.
- [Tar72] Robert E. Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972.
- [TvL84] Robert E. Tarjan and Jan van Leeuwen. Worst-case analysis of set union algorithms. *J. ACM*, 31(2):245–281, 1984.
- [Vis83] Uzi Vishkin. Implementation of simultaneous memory address access in models that forbid it. *J. Algorithms*, 4(1):45–50, 1983.