



Provably-Efficient and Internally-Deterministic Parallel Union-Find

Alexander Fedorov

Institute of Science and Technology Austria
Klosterneuburg, Austria
alexander.fedorov@ist.ac.at

Giorgi Nadiradze

Institute of Science and Technology Austria
Klosterneuburg, Austria
giorgi.nadiradze@ist.ac.at

Diba Hashemi

Institute of Science and Technology Austria
Klosterneuburg, Austria
diba.hashemi@ist.ac.at

Dan Alistarh

Institute of Science and Technology Austria
Klosterneuburg, Austria
dan.alistarh@ist.ac.at

ABSTRACT

Determining the degree of inherent parallelism in classical sequential algorithms and leveraging it for fast parallel execution is a key topic in parallel computing, and detailed analyses are known for a wide range of classical algorithms. In this paper, we perform the first such analysis for the fundamental *Union-Find* problem, in which we are given a graph as a sequence of edges, and must maintain its connectivity structure under edge additions. We prove that classic sequential algorithms for this problem are well-parallelizable under reasonable assumptions, addressing a conjecture by [Blelloch, 2017]. More precisely, we show via a new potential argument that, under uniform random edge ordering, parallel union-find operations are unlikely to interfere: T concurrent threads processing the graph in parallel will encounter memory contention $O(T^2 \cdot \log |V| \cdot \log |E|)$ times in expectation, where $|E|$ and $|V|$ are the number of edges and nodes in the graph, respectively. We leverage this result to design a new parallel Union-Find algorithm that is both *internally deterministic*, i.e., its results are guaranteed to match those of a sequential execution, but also *work-efficient and scalable*, as long as the number of threads T is $O(|E|^{\frac{1}{3}-\epsilon})$, for an arbitrarily small constant $\epsilon > 0$, which holds for most large real-world graphs. We present lower bounds which show that our analysis is close to optimal, and experimental results suggesting that the performance cost of internal determinism is limited.

CCS CONCEPTS

• **Theory of computation** → **Parallel algorithms; Graph algorithms analysis.**

KEYWORDS

union-find; parallel algorithms; graph algorithms; deterministic parallelism

ACM Reference Format:

Alexander Fedorov, Diba Hashemi, Giorgi Nadiradze, and Dan Alistarh. 2023. Provably-Efficient and Internally-Deterministic Parallel Union-Find.



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

SPAA '23, June 17–19, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9545-8/23/06.

<https://doi.org/10.1145/3558481.3591082>

In *Proceedings of the 35th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '23)*, June 17–19, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3558481.3591082>

1 INTRODUCTION

A popular approach to efficient parallelization has been to leverage the *inherent parallelism* present in many sequential algorithms, and graph problems have been shown to be a particularly fertile ground for this approach. A simple illustration is given by the greedy algorithm for Maximal Independent Set (MIS) on graphs in which we initially fix a random ordering of the graph nodes, after which we process nodes one-at-a-time, in this order, adding the node to the MIS if none of its earlier neighbors has been previously added to the MIS, and discarding it otherwise. When trying to parallelize such sequential algorithms, two key questions are 1) the likelihood of *contention*, i.e. that two arbitrary nodes in the ordering have data dependencies, and 2) the depth of the *longest dependency chain* between “dependent” nodes, for any given graph. Intuitively, the total number of nodes divided by the maximal dependency depth gives an estimate for how many nodes can be processed in parallel, on average. Assuming low dependency depth, e.g., sublinear in the number of nodes or edges, a second challenge is to design parallel algorithms which are able to leverage it for fast end-to-end runtimes, usually measured as the total number of thread memory accesses or *work*, which includes managing any auxiliary data structures.

A rich line of work has investigated these questions for many standard problems, such as MIS [5, 12], two-dimensional linear programming [7], random permutation, list, and tree contraction [20], or Delaunay mesh triangulation [7]. For example, the (expected) dependency depth for the above greedy MIS algorithm was first shown to be $O(\log^2 |V|)$ for any input graph with $|V|$ vertices [5], which was later reduced to a tight $O(\log |V|)$ via a technical breakthrough by Fischer and Noever [12]. These analyses can lead to optimal-work algorithms, and are complemented by frameworks which allow speedups to be realized in practice [8].

One classic problem for which the potential for parallelization is still not well-understood is Disjoint-Set Union / Union-Find [22], in which we are given a sequence of graph edges, and must maintain connectivity structure under edge additions. Specifically, assuming a random edge ordering and an algorithm implementing one of the standard sequential linking strategies, e.g. linking by rank [22], we wish to analyze the likelihood that processing two edges in parallel

may lead to a potential data race in the Union-Find data structure, and whether work-efficient parallel algorithms with deterministic computations exist for this problem. While Union-Find was conjectured to be provably highly-parallel by Blelloch [6], this question has so far remained open.

Contribution. In this paper, we take a significant step towards resolving this problem for natural parallel variants of the Union-Find algorithm. We start with a simple and general “edge collision” model, which captures the possible data races that may arise when executing such algorithms. We then present a new potential argument showing that, under random edge ordering, the number of edge collisions among T parallel threads can be bounded by $O(T^2 \log |V| \log |E|)$ in expectation. This implies an upper bound of expected $O(|E|^{\frac{2}{3}} \cdot \log^{\frac{1}{3}} |E| \cdot \log^{\frac{1}{3}} |V|)$ on the dependency depth of this problem. Conversely, we show an iteration lower bound of $\Omega\left(\frac{\sqrt{|E|}}{\sqrt{\log |E|}}\right)$ for cycle graphs w.h.p. for any parallel algorithm which processes maximally-large edge batches without violating dependencies, which was suggested in [7] as a general approach for randomized incremental algorithms in this setting.

Based on this new analysis, we propose a simple parallel algorithm with the total number of steps (work complexity) $O(|E|\alpha(|V|) + T^3\alpha(|V|) \log |V| \log |E|)$ in expectation, on any graph with $|E|$ edges and $|V|$ nodes, using T concurrent processors. (Here, $\alpha(\cdot)$ denotes the inverse-Ackermann function arising in the analysis of sequential Union-Find.) Thus, the overhead of parallelism, given by the second term, is negligible as long as $T = O(|E|^{\frac{1}{3}-\epsilon})$, for any constant $\epsilon > 0$, which should be reasonable in practice, as graph inputs tend to be large relative to parallelism.

Our main technical contribution is in the potential argument bounding the collision probability between tasks processing two distinct edges in the random ordering. Specifically, we assume a sequential algorithm maintaining a standard “compressed forest” data structure, where elements are nodes, arranged into directed trees, and the tree root is the set representative [22]. The addition of a new edge may lead to components being linked, where the link direction is decided by the algorithm. Two edges *collide* if they cannot be processed in parallel: for instance, this can happen when the two edges processed simultaneously would lead to the same root being linked to *different* components.

In this context, we first analyze a “sequentialized” variant of the process, in which, at each step, a new edge is added to the data structure, and consider the probability p_t of two randomly chosen edges “colliding” after exactly t steps. While we cannot bound p_t independently of the graph structure, and p_t may fluctuate significantly over steps, we are able to bound the *average* value of p_t , taken over time steps t , via a new potential function, which we show to be well-correlated with the “smoothed” collision probabilities over time, $\sum_{t=0}^{|E|-1} p_t$. Specifically, we show that, for any linking strategy that bounds the maximum forest depth to $O(\log |V|)$, the sum of collision probabilities over all edges will satisfy $\sum_{t=0}^{|E|-1} p_t = O(\log |V| \log |E|)$. In turn, this implies a bound on the expected number of collisions over a number of parallel steps. We present a complementary lower bound showing that the

number of collisions for any Union-Find algorithm is $\Omega(\log |V|)$ for cycle graphs.

Next, we design a work-efficient algorithm which leverages these analytical observations. We apply the *deterministic reservations* approach [4], which we customize to our setting. Specifically, for a well-chosen parameter S , the algorithm proceeds in “windows” of S consecutive edges, where in each such stage the threads first attempt to “mark” roots in parallel via deterministic reservations, and then examine whether any of the reservations resulted in data conflicts because of the underlying graph structure. If no such conflict occurs, then the edges can be processed fully in parallel. Otherwise, the threads process the conflict-free prefix. Then, they proceed to execute a new window of size S on the remaining suffix. The collision bound above implies that this algorithm has asymptotically-optimal work if the number of threads is $O(|E|^{\frac{1}{3}-\epsilon})$.

Our algorithm is *internally deterministic* – roughly, given a fixed input ordering, one obtains a unique dependency graph between the tasks corresponding to the edges, and will therefore have the same complexity as an equivalent sequential execution. As a consequence of internal determinism, our algorithm should be generally-useful as an efficient sub-procedure: for example, we can leverage it for a new solution for the Dynamic Spanning Tree and the Minimal Spanning Tree (MST) problems. Our main result is as follows:

THEOREM 1.1. *There exists an internally-deterministic parallel Union-Find algorithm for CRCW PRAM model with priority writes that has $O(|E|\alpha(|V|))$ expected total work and $O(|E|^{\frac{2}{3}} \cdot \text{polylog}(|E|))$ parallel depth on a randomly shuffled sequence of edges and for the number of parallel threads $T = O(|E|^{\frac{1}{3}-\epsilon})$. This also implies a parallel version of Kruskal’s MST algorithm, which has the same work and depth bounds in the special case where we have a sorted sequence of edges, and edge weights are generated so that the sorted sequence is a random shuffle.*

The full version of this paper is available in [11].

Related Work. Our work extends the line of research analyzing inherent parallelism in classical sequential algorithms [5, 7, 8, 12, 20] to show that Union-Find is also efficiently parallelizable. This partly addresses a question by Blelloch [6], who posed the dependency depth of Union-Find as an open problem.

Prior work proposed a linear-work algorithm for graph connectivity [19]; however, this uses a different decomposition-based parallelization approach, which requires a static graph, and does not allow for incremental edge additions, nor online connectivity queries. The best known parallel algorithm for the a variant of Union-Find in the concurrent-read concurrent-write (CRCW) PRAM model was proposed by Simsiri et al. [21]. They investigate a parallel version of *batched* union-find, assuming that operations are inherently grouped into batches. The algorithm is work-efficient and guarantees polylogarithmic span. It processes batches of find operations with *path compaction* and asymptotically the same total work as in the sequential setting; however, this part requires significant synchronization between threads, and thus, it may not be very efficient in practice. For *union* operations, the approach is to reduce the problem to a linear-work parallel connected components algorithm, such as [19]. By comparison, our algorithm is internally

deterministic, which allows it to be used as a sub-component for e.g. MST algorithms.

Anderson et al. [3] considered incrementally maintaining a spanning tree in a “sliding window” model, where edges may appear and disappear over time, which is different from ours. Our results should also extend to analyzing alternative ways of parallelizing sequential iterative algorithms, such as by defining task priorities and executing them via a relaxed priority queue [1, 18], in which case it can upper bound the amount of wasted work due to out-of-order execution; we leave this analysis for future work.

We also mention existing work on efficient *concurrent* variants of Union-Find algorithms, which modify the original linking approaches to employ atomic operations allowing concurrent access to the Union-Find data structure [2, 14, 15]. Our work is only partly related, as it considers a completely different parallelization approach, with different metrics and progress guarantees. Specifically, the above line of work considers an asynchronous shared-memory model with atomic operations, studying total step complexity. By contrast, we consider a standard parallel model, and study classic notions of task parallelism such as dependency depth and collisions.

2 NOTATION AND PRELIMINARIES

Arguably, the most common data structure to address the Disjoint-Set Union / Union-Find problem is the *compressed forest* [22]. Here, nodes / set elements form directed trees with roots used as *representatives* of their sets. As a result, checking whether two elements are in the same set can be implemented by following *parent links* of these elements in the directed trees, finding the root representatives, and comparing them. Merging two sets means adding a link from one root to another one.

Algorithms based on *compressed forest* data structures differ in two key ingredients. The first is the *linking technique*: by choosing which root becomes the common root when merging two trees, it is possible to limit tree depth. Three popular linking strategies are *linking by size*—linking the smaller tree to the larger one, *linking by rank*—similar to *linking by size* but tree sizes are approximated, and *linking by random priorities*—the root with lower random priority is linked to the root with higher priority. All three linking strategies achieve $O(\log n)$ maximum tree depth [13, 22], though in the case of *linking by random priorities* this is in expectation. The second key component is *path compaction*—traversed paths in trees are shortened after every root search by replacing parents of every visited node with nodes higher in the tree. When combined with any of the linking techniques, this results in $O(\alpha(n))$ amortized time per operation, where $\alpha(n)$ is the inverse-Ackermann function.

We will associate $\text{unite}(u, v)$ operations with (u, v) edges in the union graph $G = (V, E)$, where vertices correspond to Union-Find elements. These edges can also be viewed as tasks; to execute the task means to unite sets corresponding to edge endpoints. This high-level approach allows us to study union-find properties for different graph structures.

2.1 Definition of Collisions

The key definition for analysing parallel and concurrent dynamic graph algorithms is that of an *edge collision*. Intuitively, two edges *collide* when they cannot be processed in parallel. To define what

this means for the Union-Find problem, consider Listing 1. When joining sets (i.e., trees), the classical Union-Find implementation compares their sizes and then links the root of the smaller one to the root of the larger tree (lines 6 and 9). It is easy to see that concurrent links of the same root to different trees result in the loss of a link, and thus, are incorrect. In practice, this is where concurrent Union-Find algorithms employ synchronization primitives (e.g., Compare-And-Set in [14]). We say that two edges *collide* when they both connect different components and share the “smaller” component (the exact order of the components is determined by the linking strategy). Note that size updates in lines 7 and 10 are much less crucial, since they commute and can be implemented with simple atomic operations such as `fetch-and-add`.

```

1 bool union(u, v):
2   u := find_root(u)
3   v := find_root(v)
4   if u == v: return false // Already in the same component
5   if u.size < v.size: // Which component is larger?
6     u.parent = v // Link u to v
7     v.size += u.size
8   else:
9     v.parent = u // Link v to u
10    u.size += v.size
11  return true

```

Listing 1: union operation in the Union-Find data structure with the *linking by size* technique.

One may ask whether the definition of collision can be simplified to requiring that “two edges share an adjacent tree.” This change will not make data structure implementations simpler but will simplify the analysis. Unfortunately, this dramatically increases the number of collisions. Consider Erdos-Rényi random graphs [10] with the number of already inserted edges between $(\frac{1}{2} + \epsilon)|V|$ and $(1 - \epsilon)|V|$, where ϵ is a constant strictly between 0 and $\frac{1}{4}$. It is known that the largest component in this case is of size $\Omega(|V|^{2/3})$, and all other components are of sizes $O(\log |V|)$, w.h.p. [10]. As a result, the probability of *simplified* collision for two random edges is $\Omega(|V|^{-1/3})$, while in the case of our definition it is $O\left(\frac{\log |V|}{|V|}\right)$. In fact, this suggests that for our definition the total number of collisions for random graphs is polylogarithmic, whereas for the simplified definition it is polynomial.




Figure 1: An example of two edges colliding in the case of linking by size. The black solid edges are the ones already inserted. The dashed edges are yet to be added to the data structure. Since C_3 is the smaller component, the blue edges will try to write to the same memory location (the root of C_3), i.e., they have a collision and cannot be added to the data structure in parallel.

2.2 The Random Process

Similar to prior work, e.g. [4, 5], it is convenient to assume that the operations of the Union-Find algorithm are executed sequentially, i.e., edges are added one by one to the union graph G in some order. This may not be the case in the actual execution, since parallelism allows to add several edges at once, although, in our parallel algorithm, to guarantee deterministic computations, these edges will still be from the prefix of the unprocessed edge sequence. We call “step t ” the moment of time when the t -th edge has been added to the data structure. Specifically, we denote the t -th added edge by ε_t and the current set of inserted edges is $E_t = \{\varepsilon_1, \dots, \varepsilon_t\}$. Edges that are yet to be added (i.e., edges from $E \setminus E_t$) are called *active*. In each step, we assume in our analysis that all present connected components are enumerated in the order defined by the used linking strategy (e.g., in order of increasing sizes when *linking by size* is used) from 1 to C_t . Finally, we call m_i^t the number of active edges that connect the component of index i to a connected component with a strictly greater index (i.e., to a “larger” component) in step t . Notice that these m_i^t edges pairwise collide according to our definition of collision.

As in prior works analysing graph algorithms [5, 7, 20], we assume that the order of edges is uniformly random. Otherwise, there can be examples with $\Omega(|E|)$ collisions among subsequent pairs of edges in the sequence, resulting in no potential for parallelism. We will prove that the expected number of collisions among two random *active* edges is small. More formally, let X_t be an indicator random variable for the event “two distinct active edges (i.e. from $E \setminus E_t$) chosen uniformly at random in step t collide.” Let the probability of collision for two random distinct active edges in step t be $p_t = \mathbb{E}[X_t]$. We will use conditional expectation $\mathbb{E}[X_t \mid \varepsilon_1, \dots, \varepsilon_t]$ when previously added edges are known (ε_t is a random variable when the order of edges is random).

The next theorem is the core of our performance analysis in Section 4, and is proven in Section 3. It states that the sum of probabilities of collisions over all steps is polylogarithmic.

THEOREM 3.7. *For a uniform random edge ordering, and any linking Union-Find strategy that bounds union-find forest depth to $\sum_{t=0}^{|E|-1} p_t = O(\log |V| \cdot \log |E|)$, where $|V|$ is the number of vertices and $|E|$ is the number of edges.*

2.3 Model of Parallelism

For our parallel algorithm (Subsection 4.2), we assume concurrent-reads concurrent-write parallel random-access machine (CRCW PRAM) with *priority write*. Priority write (`w r i t e _ m i n`) is an atomic instruction that writes an input value to some memory address only if it is smaller than the current value at that location. This atomic instruction is widely used in parallel algorithms, especially for the Connected Components and Spanning Tree problems, and can be efficiently implemented [7, 19].

We analyse our parallel algorithms using the standard work-depth approach. The *work* of a parallel algorithm is the total number of executed instructions. The *parallel depth* (or *span*) is the length of the longest computational dependency chain.

3 COLLISION ANALYSIS

The goal of this section is to introduce an argument for bounding the number of collisions in the case of random edge ordering. First of all, we will prove the following lemma about collision probability when all added edges are known. Please recall that all notation has been defined in Subsection 2.2.

LEMMA 3.1. *For any $t \geq 0$, the collision probability in step t for fixed $\varepsilon_1, \dots, \varepsilon_t$ is $\mathbb{E}[X_t \mid \varepsilon_1, \dots, \varepsilon_t] = \sum_{i=1}^{C_t} \frac{m_i^t}{|E|-t} \cdot \frac{m_i^t-1}{|E|-t-1}$.*

PROOF. The probability that we select an edge from component i to a larger component is the number of such edges, denoted by m_i^t , divided by the number of remaining edges in step t , which is $|E| - t$. The only edges that can cause a collision with the chosen edge are other edges connecting component i to a larger component. Consequently, the probability that another random edge causes a collision with regard to selected edge is $\frac{m_i^t-1}{|E|-t-1}$, and therefore, the probability that two random edges collide is $\sum_{i=1}^{C_t} \frac{m_i^t}{|E|-t} \cdot \frac{m_i^t-1}{|E|-t-1}$. \square

Overview. A natural proof strategy would be to directly bound p_t . However, the value of p_t can heavily depend on the structure of graph G . For example, it follows by simple calculation that for a star graph p_0 is at least some constant greater than zero, while for a cycle graph $p_0 = o(1)$, but $p_{|E|-3}$ is a non-zero constant. Instead, we prove that p_t cannot be high for all possible t by employing a type of amortized analysis. Specifically, we present a metric which increases considerably when p_t is high, and is bounded at the same time. As a result, this will help us to bound $\sum_{t=0}^{|E|-1} p_t$.

Let $\text{depth}_t(u)$ be the depth of vertex u in the *uncompressed* Union-Find forest in step t , where “uncompressed” means that depth is counted ignoring any previous path compaction. Most reasonable linking techniques guarantee that for any vertex u , we have that $\text{depth}_t(u) \leq C \cdot \log |V|$ for some constant $C > 0$. We define our potential function as follows:

Definition 3.2. For each edge $(u, v) \in E$, we define its rank at time t as $\phi_t(u, v) = \text{depth}_t(u) + \text{depth}_t(v)$. The potential function at t will be $\Phi_t = \frac{|E|}{|E|-t} \sum_{e \in E \setminus E_t} \phi_t(e) + \sum_{i=1}^t \frac{|E|}{|E|-i} \cdot \phi_i(\varepsilon_i)$, $0 \leq t \leq |E| - 1$.

In more detail, the potential function is a sum over all edges of their ranks, but the ranks are multiplied by a constantly growing factor $\frac{|E|}{|E|-t}$, and the rank ϕ_i and its multiplier for edge ε_i are “frozen” at time i when the edge was processed.

Initially, we have $\Phi_0 = 0$, as the graph is empty, and thus, the ranks of all edges are zero. Let us observe how this potential is expected to change after each step, given that all remaining edges have equal probability to be added to the data structure. That is, we follow the sum $\mathbb{E}[\Phi_1 - \Phi_0] + \mathbb{E}[\Phi_2 - \Phi_1 \mid \varepsilon_1] + \dots + \mathbb{E}[\Phi_{|E|-1} - \Phi_{|E|-2} \mid \varepsilon_1, \varepsilon_2, \dots, \varepsilon_{|E|-2}]$. Note that, by the law of total expectation, $\mathbb{E}[\Phi_1 - \Phi_0] + \dots + \mathbb{E}[\Phi_{|E|-1} - \Phi_{|E|-2} \mid \varepsilon_1, \varepsilon_2, \dots, \varepsilon_{|E|-2}] = \sum_{t=1}^{|E|-1} \mathbb{E}[\Phi_t - \Phi_{t-1}] = \mathbb{E}[\Phi_{|E|-1}]$. Next, we connect this potential with the probability of collision.

$$\text{LEMMA 3.3. } \mathbb{E} [\Phi_{t+1} - \Phi_t \mid \varepsilon_1, \dots, \varepsilon_t] \geq |E| \cdot \sum_{i=1}^{C_t} \frac{m_i^t}{|E|-t} \cdot \frac{m_i^t}{|E|-t-1}.$$

PROOF. Assume that in step $t + 1$, the newly added edge ε_{t+1} is $(u, v) \in E$. If the edge is internal, i.e., connects nodes in the same component, the edge does not affect connectivity or edge ranks. The potential function still increases because of the increasing multiplier $\frac{|E|}{|E|-t}$, but we use a trivial lower bound of 0 for this increase. If the edge is external, i.e., connects two different connected components, assume that the index of the smallest component is i . Then, the number of edges connecting this component to larger components is m_i^t , so the probability that this component is chosen as the smaller one is $\frac{m_i^t}{|E|-t}$. Due to the fact that the root of this component will be linked to another component, ranks of all active edges adjacent to it are increased by at least 1. We can lower bound the number of such edges by m_i^t . This means that, with probability at least $\frac{m_i^t}{|E|-t}$, component i is chosen and then the sum of active edge ranks is increased by at least m_i^t . Finally, taking into consideration the multiplier of the potential function and summing over all components, we get the required inequality

$$\begin{aligned} \mathbb{E} [\Phi_{t+1} - \Phi_t \mid \varepsilon_1 \dots \varepsilon_t] &\geq \frac{|E|}{|E|-t-1} \cdot \sum_{i=1}^{C_t} m_i^t \cdot \frac{m_i^t}{|E|-t} = \\ &= |E| \cdot \sum_{i=1}^{C_t} \frac{m_i^t}{|E|-t} \cdot \frac{m_i^t}{|E|-t-1}. \end{aligned}$$

□

LEMMA 3.4. *The increment of the potential at each step $t \geq 0$ satisfies $\mathbb{E} [\Phi_{t+1} - \Phi_t \mid \varepsilon_1, \dots, \varepsilon_t] \geq |E| \cdot \mathbb{E} [X_t \mid \varepsilon_1, \dots, \varepsilon_t]$.*

PROOF. Follows by combining Lemma 3.1 with the inequality in Lemma 3.3. Note that these expectations are different – the left one is over the choice of ε_{t+1} , while the right one is over the choice of two random active edges. □

$$\text{COROLLARY 3.5. } \mathbb{E} [\Phi_{|E|-1}] \geq |E| \cdot \sum_{t=0}^{|E|-2} p_t.$$

PROOF. Summing up inequality from Lemma 3.4 over all steps and taking expectation of the sum, we get the following:

$$\mathbb{E} \left[\sum_{t=0}^{|E|-2} \mathbb{E} [\Phi_{t+1} - \Phi_t \mid \varepsilon_1, \dots, \varepsilon_t] \right] \geq |E| \cdot \mathbb{E} \left[\sum_{t=0}^{|E|-2} X_t \mid \varepsilon_1, \dots, \varepsilon_t \right]$$

By the law of total expectation, this can be simplified to the next inequality:

$$\mathbb{E} \left[\sum_{t=0}^{|E|-2} \Phi_{t+1} - \Phi_t \right] \geq |E| \cdot \mathbb{E} \left[\sum_{t=0}^{|E|-2} X_t \right]$$

Finally, using the definition of p_t and reducing the left sum, we deduce that:

$$\mathbb{E} [\Phi_{|E|-1}] = \mathbb{E} [\Phi_{|E|-1} - \Phi_0] \geq |E| \cdot \sum_{t=0}^{|E|-2} p_t.$$

□

This corollary means that in order to upper bound the smoothed collision probabilities $\sum_t p_t$, we can bound the potential function instead. We will do this in the next lemma. Note that for randomized linking strategies (e.g., *linking by random priorities*) the next bound holds in expectation over the randomness used in linking.

LEMMA 3.6. *For any linking strategy bounding maximal depth to $\mathcal{O}(\log |V|)$, we have $\Phi_{|E|-1} = \mathcal{O}(|E| \cdot \log |V| \cdot \log |E|)$.*

PROOF. When forest depth is $\mathcal{O}(\log |V|)$, it is easy to see that every edge rank is $\mathcal{O}(\log |V|)$, and thus, the sum of edge ranks is $\mathcal{O}(|E| \cdot \log |V|)$. The problem is the $\frac{|E|}{|E|-t}$ edge rank multiplier, which ranges from 1 to $|E|$. That is why we “frozen” ranks of edges upon their processing.

Observe the edge rank multipliers for edges added in different steps. For the first $\frac{|E|}{2}$ edges, this multiplier is at most $\frac{|E|}{|E|-|E|/2} \leq 2$. Similarly, for the next $\frac{|E|}{4}$ edges, this multiplier is at most $\frac{|E|}{|E|-3|E|/4} \leq 4$. By continuing this progression, we get the following for forest depth $\leq C \cdot \log |V|$:

$$\begin{aligned} \Phi_{|E|-1} &\leq \frac{|E|}{2} \cdot 2 \cdot C \cdot \log |V| \cdot 2 + \frac{|E|}{4} \cdot 4 \cdot C \cdot \log |V| \cdot 2 + \dots = \\ &= \sum_{i=0}^{\log_2 |E|} |E| \cdot C \cdot \log |V| \cdot 2 = \mathcal{O}(|E| \cdot \log |V| \cdot \log |E|). \end{aligned}$$

□

THEOREM 3.7. *The sum of collision probabilities satisfies $\sum_{t=0}^{|E|-1} p_t = \mathcal{O}(\log |V| \cdot \log |E|)$ given a random order of edges and any linking Union-Find strategy that bounds forest depth to $\mathcal{O}(\log |V|)$.*

PROOF. Combining Corollary 3.5 and Lemma 3.6, we conclude $\sum_{t=0}^{|E|-2} p_t \leq \frac{1}{|E|} \cdot \Phi_{|E|-1} = \mathcal{O}(\log |V| \cdot \log |E|)$. □

Surprisingly, the proof of Theorem 3.7 shows $\sum_{t=0}^{|E|/2} p_t \leq C \cdot \log |V|$, i.e., for the first half of edges the expected number of collisions is just $\mathcal{O}(\log |V|)$, which may mean that the factor of $\log |E|$ is an artefact of the current analysis.

4 WORK-EFFICIENT UNION-FIND

We start by showing how to apply Theorem 3.7 to some practical Union-Find algorithms in various computation models. Specifically, we will speak about incremental dynamic connectivity and minimum spanning tree problems.

4.1 Bounding Contention

The discussion so far assumed an idealized sequential execution. We now wish to analyse concurrent Union-Find algorithms, which can be implemented in practice via locks, hardware transactions, or lock-free primitives [2, 9], using a generalized concurrency model. In all these implementations, extra work comes from *write contention*—situation when several threads try to modify the same memory location simultaneously. For locks, write contention means the need of waiting until another thread finishes modifying the required

memory location, while for hardware transactions and lock-free primitives, it means retries of operations. Moreover, in practice, write contention causes additional L3 cache misses in NUMA (Non-Uniform Memory Access) computer architectures. This is why we will focus on bounding the probability that threads experience write contention when attempting to process edges in parallel.

More formally, our basic concurrency model is: a set of $T \geq 2$ threads execute synchronously in iterations, where in each iteration each thread picks a remaining edge uniformly at random, and inserts it into the Union-Find data structure. The threads may pick colliding edges, in which case we register a *write contention* event; otherwise, the edges are processed without contention. Either way, we assume that the edges are processed. Our goal is to bound the expected total number of write contention events between threads when processing all edges. In our case, write contention between threads means collision between their edges according to our definition of collision.

THEOREM 4.1. *The expected number of write contention events for two threads in this concurrent model is bounded by $O(\log |V| \cdot \log |E|)$.*

PROOF. Since edges are inserted randomly into the data structure, this algorithm is almost the same as the random process described in Subsection 2.2. Specifically, iteration t in this algorithm corresponds to step $2t$ in the random process because two new edges are added every iteration. Moreover, the probability of write contention at iteration t is exactly p_{2t} (see the definitions in Subsection 2.2). So, the total number of write contention events is $\sum_{t=0}^{|E|/2} p_{2t} \leq \sum_{t=0}^{|E|} p_t = O(\log |V| \cdot \log |E|)$. The last inequality is from Theorem 3.7, stated in the end of the previous Section. \square

THEOREM 4.2. *The expected number of write contention events for T threads in this concurrent model is $O(T^2 \cdot \log |V| \cdot \log |E|)$.*

PROOF. The number of possible pairwise contentions every iteration is $\Theta(T^2)$, so by linearity of expectation, the expected number of collisions at iteration t is $O(T^2 \cdot p_{T \cdot t})$. Then, the total number of contention events is bounded by $\sum_{t=0}^{|E|/T} T^2 \cdot p_{T \cdot t} \leq \sum_{t=0}^{|E|} T^2 \cdot p_t = O(T^2 \cdot \log |V| \cdot \log |E|)$. \square

The last two results show that the number of write contention events depends on the size of the graph only polylogarithmically, and thus, may be dominated by other costs of processing a graph. This suggests fairly low contention cost for implementing Union-Find via most concurrency primitives (locks, atomic operations, transactions), which was also previously shown in practice [9].

4.2 Parallel Algorithm

Our main goal is to design a parallel iterative algorithm for Union-Find that is work-efficient and internally deterministic. While work-efficient parallel algorithms are well-known for many problems, relatively few algorithms have the second property [4, 7, 20]. Internal determinism is particularly interesting for Union-Find, since running an internally-deterministic parallel Union-Find algorithm on a sorted sequence of edges is similar to Kruskal’s algorithm [17] and results in a minimal spanning tree. Blelloch et al. [4] present a

practical algorithm for Union-Find but without a theoretical analysis. Their algorithm is deterministic but *not internally deterministic*, as its output (i.e., edges used for connectivity) depends on some predefined parameter and may diverge from the output of the sequential algorithm.

We aim to close these gaps by providing a practically-efficient algorithm which is provably work-efficient, and *internally deterministic*. Specifically, the algorithm uses concurrent-reads concurrent-write (CRCW) PRAM with *priority write* (see Subsection 2.3).

Deterministic Reservations. We build on the algorithm of Blelloch et al. [4], which uses the *deterministic reservations* approach. Specifically, in each iteration, their algorithm considers the prefix of remaining tasks (i.e. edges) of size S and then proceeds in two phases. In the first phase, these tasks do *reservations* of memory locations they want to change by using *priority write*. In the second phase, all tasks in the prefix that succeeded in their reservations are executed. For the Union-Find problem, in the first phase tasks reserve the smaller adjacent connected component for each edge, by making a priority write in its root and then link roots of these smaller components to other components in the case of successful reservation. This behaviour follows our definition of collision: the algorithm adds all edges in the prefix that do not collide with any preceding edges.

There are several challenges when making this algorithm both work-efficient and internally deterministic. First, Blelloch et al. do not analyze how often collisions happen and how many iterations the algorithm requires. Furthermore, parallelism in their algorithm violates linking strategy properties, which may make their algorithm not work-efficient. For example, on a directed path, their algorithm can add all edges in one iteration, but the depth of the resulting Union-Find forest will be linear. Last but not least, the algorithm is not internally deterministic: there are cases where the sequential algorithm uses some edge for connectivity (i.e., $\text{union}(u, v)$ returns true), but this algorithm replaces it with an edge located later in the sequence. Blelloch et al. suggest a modification of their algorithm to solve Spanning Tree, which does make it internally deterministic but it requires each edge to reserve *both* of its endpoints. As a result, for random or star graphs, few operations will succeed in every iteration, since most operations will try to reserve the largest connected component.




Figure 2: An example how our parallel algorithm processes edges. It scans some prefix of the remaining edge sequence and does deterministic reservations. If there is a collision, i.e., for some edge, deterministic reservation failed because of another edge, the algorithm processes not all considered prefix, but only its part before the first collision.

Internally-Deterministic Algorithm. One key difference between the algorithm of Blelloch et al. and the one we are analyzing is that we do not allow the execution of a task unless all tasks before it will be executed by the end of the current iteration (Figure 2). In other words, we execute all tasks in the prefix until a task with failed reservation, i.e., an edge with a collision with one of the previous edges. This modification addresses the latter problem in the algorithm of Blelloch et al. by making it internally deterministic. This appears to make the algorithm less efficient; yet, we will prove that collisions are relatively rare, so the modification does not have significant impact on performance. Similarly to other iterative algorithms [7, 20], we assume that the order of edges is uniformly random.

```

1 void make_reservations(edges, l, r):
2   parallel for t ∈ [l, r):
3     (u, v) := edges[t] // alias for edge endpoints
4     u_root = find_root(u)
5     v_root = find_root(v)
6     if u_root == v_root:
7       continue // already in the same component
8     if u_root < v_root: // linking strategy
9       write_min(u_root.reservation, t)
10    else:
11      write_min(v_root.reservation, t)
12
13 void union_find(edges):
14   i := 0
15   while i < |edges|:
16     make_reservations(edges, i, i + S):
17     stop := min edge id with failed reservation or i + S
18     sz := stop - i // max size without collisions
19     Parallel-Link-All(i, i + sz)
20     i += sz

```

Listing 2: Parallel algorithm for Union-Find. In each iteration, a prefix of size S is being processed in parallel unless there is a write collision among its tasks. Parallel-Link-All was presented in the work-efficient Union-Find [21]. It groups roots according to connected components and then unites them via a recursive divide-and-conquer strategy.

Our parallel algorithm is presented in Listing 2. It starts the same way as the algorithm of Blelloch et al. [4] by making *deterministic reservations* in the unprocessed prefix of size S (Line 16). Specifically, the algorithm finds the roots of the corresponding trees for each edge and tries to reserve the smaller one according to the used linking strategy. Then, it checks in parallel whether all reservations are successful and, if not, it finds the first unsuccessful reservation using `write_min` instruction (Line 17). Finally, the algorithm completes all edge additions in the prefix until the first failed reservation (Line 19). In the simplest case, this means just linking the root that was designated as “smaller” in the previous step to the larger one. However, we will propose another approach that maintains all the guarantees of the linking techniques on forest depth. This process is repeated until all edges are processed.

Optimizations. A naive implementation would spend $O(S \cdot \text{depth})$ total work and has $O(\text{depth})$ span for each iteration, where *depth* is the Union-Find forest depth; due to parallelism, even with linking

techniques, the forest depth potentially can be up to linear in the simplest algorithm. However, incremental improvements can address this and improve total work to an optimal $O(S \cdot \alpha(|V|))$ while keeping the span polylogarithmic, where $\alpha(\cdot)$ is the inverse Ackermann function. Both adjustments were first proposed by Simsiri et al. [21] in their work-efficient batched parallel Union-Find algorithm. First of all, Parallel-Link-All can preserve forest depth bounds if a linear-work connected components parallel algorithm is used to group all connected components that will form one component [21]. Then, all grouped connected components can be merged in parallel by applying a recursive divide-and-conquer approach. This Parallel-Link-All takes linear work and polylogarithmic span. The second optimization ensures that parallelism in path compaction does not increase total work. Specifically, in lines 4 and 5 the same bulk-parallel approach as in [21] is employed. Roughly, the algorithm acts like a parallel BFS (Breadth-First Search) that starts in vertices, for which we want to find the root, and ascends wave-by-wave until it finds all roots. Then it traverses all nodes again and re-links them directly to the roots. These adjustments allow Union-Find to use both path compaction and linking strategies, and as a result, the average work per edge is the optimal $O(\alpha(|V|))$.

4.3 Algorithm Analysis

THEOREM 4.3. *The Union-Find algorithm in Listing 2 is internally deterministic.*

PROOF. We proceed by contradiction. Assume that the sequential algorithm uses a different set of edges to form connected components. Consider the first edge (u, v) for which the result differs. It was either taken by the presented algorithm and not taken by the sequential one or vice versa.

Case 1: This edge was used by our algorithm but not by the sequential one. The fact that it was not used by the sequential algorithm means that it, together with the previous edges, forms a cycle. As it was the first difference, our algorithm acted incorrectly and created a cycle in the Union-Find forest. However, linking techniques impose a total order on all connected components and the algorithm links smaller roots to larger ones. The cycle should have all edges directed in the same direction since in one iteration only one outgoing edge can be added for each connected component. This means that obtaining a cycle is impossible, because in a directed cycle at least one edge (link) contradicts the total order at the start of iteration and our algorithm does not allow this.

Case 2: (u, v) was used by the sequential algorithm and not by our algorithm. Consider the iteration when the edge was processed. We know that, at each iteration, all edges inside the processed prefix either tried to make a reservation and succeeded, or did not try because the vertices are already connected. We know the former is not true, because by assumption, (u, v) was not used, so the latter should be true, which means that u and v were already connected before the iteration. However, we know that the edges added before the iteration cannot connect u and v , otherwise the sequential algorithm would make a cycle. We have obtained a contradiction. \square

After proving that our algorithm is indeed internally-deterministic, what is left is to bound the number of its iterations.

THEOREM 4.4. *The expected number of iterations of the Union-Find algorithm in Listing 2 is $\left\lceil \frac{|E|}{S} \right\rceil + O(S^2 \cdot \log |V| \cdot \log |E|)$, if the order of edges is uniformly random.*

PROOF. It is easy to see that the minimum number of required iterations in absence of collisions is $\left\lceil \frac{|E|}{S} \right\rceil$, since at most S edges are processed in each iteration. Every collision in the prefix can cause the algorithm to perform at most one more iteration. So, to bound the number of “extra” iterations we can bound the total number of collisions across all considered prefixes of unprocessed edges.

Let k be an arbitrary index in the edge sequence. We want to bound the probability of collision inside the window denoted by positions $[k, k + S)$ if all edges before k were inserted in the data structure. For this reason, we define Y_k to be an indicator random variable for the event “there is at least one collision inside window $[k, k+S)$ if all edges before k were inserted in the data structure”. Since the edge sequence is random, the inserted edges and edges inside this window are random as well. For two fixed positions in the window the probability of their collision then is exactly $p_k = \mathbb{E}[\mathbb{E}[X_k \mid \varepsilon_1, \dots, \varepsilon_k]]$ (see Subsection 2.2). By the union bound, the probability of a collision in the window $\mathbb{E}[Y_k \mid \varepsilon_1, \dots, \varepsilon_k]$ is then $\leq S^2 \cdot p_k$, since there are less than S^2 pairs of possibly colliding positions in the window.

Let I be the set of unprocessed prefix start positions (i.e., the set of all observed values of i in Listing 2). Similarly to the previous paragraph, we would want to say that the total expected number is bounded by $\sum_{i \in I} S^2 \cdot p_i$. However, the content of I has a correlation with the sequence of edges. For instance, if in some iteration there was a collision, we know that the unprocessed prefix of the next iteration starts with an edge that had a collision with one of the edges before. This means that the order of edges in the unprocessed tail is not uniformly random anymore and we cannot use the definitions from Subsection 2.2.

This issue has two possible solutions. The first one requires a small modification to Listing 2: instead of processing the window $[i, stop)$, the algorithm will process this window and then will separately process the task at position $stop$. This modification does not asymptotically increase the total work or the span of an iteration. However, the separately processed edge at position $stop$ removes the correlation between the order of remaining edges and the start position of the next prefix. This is because the edges strictly after position $stop$ do not affect the current iteration, and thus, all possible permutations of the tail continue to be equiprobable.

The second solution does not require modifying Listing 2 but complicates the analysis. Consider a toy algorithm which adds edges one by one and at the same time counts the number of pairwise collisions in *all* possible windows of size S . Note that this algorithm does not have the same problem as the algorithm in Listing 2 and all permutations of the tail of edges continue to be equiprobable, since the algorithm tries *all* windows of size S . This allows us to say that the probability of a collision in window i is $\mathbb{E}[Y_i]$, and thus, by the union bound, the total expected number of collisions is $\mathbb{E}\left[\sum_{i=0}^{|E|-1} Y_i\right] \leq \sum_{i=0}^{|E|-1} S^2 \cdot p_i = O(S^2 \cdot \log |V| \cdot \log |E|)$. The last equality is Theorem 3.7. Finally, the last observation we have to make is

that the number of collisions counted by the toy algorithm is not less than the number of collisions that occurred in the algorithm in Listing 2 due to the fact that all windows considered in the latter algorithm were also considered by the former algorithm.

Both approaches show that only $\sum_{i=0}^{|E|-1} S^2 \cdot p_i = O(S^2 \cdot \log |V| \cdot \log |E|)$ expected collisions happen in the algorithm, and consequently, the number of extra iterations is $O(S^2 \cdot \log |V| \cdot \log |E|)$. \square

COROLLARY 4.5. *The proposed Union-Find algorithm is work-efficient in expectation when $S = O(|E|^{\frac{1}{3}-\varepsilon})$ for any constant $\varepsilon > 0$.*

PROOF. We know that one iteration takes $O(S \cdot \alpha(|V|))$ total work. As a result, the expected total work is the expected number of iterations multiplied by the iteration work, which is $O(S \cdot \alpha(|V|) \cdot (\frac{|E|}{S} + S^2 \cdot \log |V| \cdot \log |E|)) = O(|E| \cdot \alpha(|V|) + S^3 \cdot \log^3 |E|) = O\left(|E| \cdot \alpha(|V|) + |E| \cdot \frac{\log^3 |E|}{|E|^{3\varepsilon}}\right) = O(|E| \cdot \alpha(|V|))$. \square

COROLLARY 4.6. *The number of iterations for the presented Union-Find algorithm is $O(|E|^{\frac{2}{3}} \cdot \log^{\frac{1}{3}} |E| \cdot \log^{\frac{1}{3}} |V|)$ in expectation when S is chosen optimally. The expected parallel depth (span) of the algorithm is then $O(|E|^{\frac{2}{3}} \cdot \text{polylog}(|E|))$.*

PROOF. It is easy to prove by substitution that the asymptotic minimum of the equation in Theorem 4.4 is achieved when $S = \sqrt[3]{\frac{|E|}{\log |V| \cdot \log |E|}}$. The span directly follows from the bound on the number of iterations and the polylogarithmic span of one iteration. \square

In addition, we observe that when the prefix size is S , up to S processes can be leveraged by the algorithm. When combined with Corollary 4.5, we get the following theorem:

THEOREM 4.7. *There exists an internally-deterministic parallel Union-Find algorithm for CRCW PRAM model with priority write that has $O(|E| \cdot \alpha(|V|))$ expected total work on a randomly shuffled sequence of edges and scales perfectly for up to $O(|E|^{\frac{1}{3}-\varepsilon})$ processes.*

It is worth mentioning that for dense graphs this total work becomes linear in the number of edges. This is because the sequential Union-Find data structure has another work bound of processing $|E|$ edges, which is $O(|E| + |V| \cdot \lg^*(|V|))$. As a result, when $|E| > |V| \cdot \lg^*(|V|)$, this bound is better than $O(|E| \cdot \alpha(|V|))$ and is just $O(|E|)$.

Iteration Dependence Depth. Another metric of interest is the *iteration dependence depth* [7]. Yet, the definition of this metric is not obvious for the Union-Find problem. A natural definition could be that *there is a dependency between tasks if at some time there is a data race between them*. Following this definition, Corollary 4.6 means that the expected iteration dependency depth is $O(|E|^{\frac{2}{3}} \cdot \log^{\frac{1}{3}} |E| \cdot \log^{\frac{1}{3}} |V|)$: by construction, there are no data races in the same processed prefix, which results in the iteration depth being bounded by the number of processed prefixes.

A more general definition for dependency depth could be that exactly the same set of union(u , v) operations should succeed (i.e., return true). However, for this definition, the dependence depth




Figure 3: An example of Union-Find execution on a cycle graph. Solid lines symbolize already added edges, while dotted lines are edges that are not inserted yet. The red edge was randomly chosen, its clockwise next component is smaller than its neighbours, so the only edge this edge can have collision with it is the green edge.

is constant: failed $\text{union}(u, v)$ operations depend on successful union operations that formed the path between u and v , while successful operations do not have any dependencies.

4.4 Simple Extension to MST

Lastly, since the algorithm is internally deterministic, it works similarly to a parallel version of Kruskal’s algorithm on a sorted sequence of edges. We believe this is the first direct parallelization of Kruskal’s algorithm. For example, in the parallel MST algorithm of Katsigiannis et al. [16], the main thread processes the whole edge sequence, while other threads are helping by checking and filtering out internal edges. In this case, the work of *successful* $\text{unite}(u, v)$ operations was still essentially sequential and done by the same thread. Moreover, in Blelloch et al. MST algorithm [4], few edges are processed on average in one iteration for random and star graphs. By contrast, we can parallelize this process here. The key restriction is that, to be provably efficient, our approach requires the sorted sequence of edges to be in fact a random shuffle of edges.

COROLLARY 4.8. *There exists a MST algorithm for CRCW PRAM model with priority write that has $O(|E| \cdot \alpha(|V|))$ expected total work, when random edge weights are generated independently and the edge sequence is already sorted, and it scales for $T = O(|E|^{\frac{1}{3}-\epsilon})$ parallel threads, for any constant $\epsilon > 0$.*

5 LOWER BOUNDS

In the previous sections, we showed upper bounds for different metrics, leaving open the question of tightness of these bound. We now attempt to close this gap. Our proofs require the reasonable assumption that the Union-Find tie-breaker when there are isomorphic connected components (in particular, of the same rank or size) acts randomly, i.e., compares random priorities, since it does not have any additional information about the graph structure.

5.1 Lower Bound on Collision Count

We start by providing a lower bound on the number of collisions (Theorem 3.7). Specifically, we show that the poly-logarithmic bound we provided cannot be improved by more than a logarithmic factor.

THEOREM 5.1. *For a cycle graph, we have $\sum_{t=0}^{|E|-1} p_t = \Omega(\log |V|)$.*

PROOF. Consider a cycle graph. In step t , there are exactly $|V| - t$ connected components and $|V| - t$ active edges. Note that since we added edges in a uniformly random order, any combination of t edges from the cycle is equally probable to be picked. For every connected component, there are two neighboring components (i.e., components adjacent through an active edge) in any step t , when $t \leq |V| - 3$. An important observation is that if we contract every currently present connected component in a single vertex, the remaining graph is still a cycle. When we add a random edge, for the clockwise-next endpoint component there is a $\frac{1}{3}$ chance that it is smaller than both of its neighbouring components due to symmetry of the graph (unless there are two or less components in total, but this occurs only in the last 3 steps). If it is smaller than its neighbours, then there is exactly one edge (out of $|V| - t$ edges) in the graph that can cause a collision with the chosen one – the edge between this endpoint and its other neighbouring component (Figure 3). So, we get:

$$\sum_{t=0}^{|V|-1} p_t \geq \sum_{t=0}^{|V|-3} p_t \geq \sum_{t=0}^{|V|-3} \frac{1}{3} \cdot \frac{1}{|V| - t - 1} = \Theta(\log |V|).$$

□

This implies the following for the model in Section 4.1:

THEOREM 5.2. *The expected write contention for two threads in the concurrent model is $\Omega(\log |V|)$ for cycle graphs.*

PROOF. As we know from Theorem 5.1, for a cycle graph $p_t \geq \frac{1}{3} \cdot \frac{1}{|V| - t - 1}$. On the other hand, from Theorem 4.1, the expected number of write contentions is $\sum_{t=0}^{|E|/2} p_{2t}$. As a result, the required bound is the sum of harmonic series over every second element and is still $\Theta(\log |V|)$. □

5.2 Lower Bound on Number of Iterations

A significantly more involved analysis is required for the parallel algorithm from Section 4.2. The challenge is that we cannot use the symmetry argument to show that the probability of collision for an adjacent pair of edges is $\frac{1}{3}$, since the subsequences of edges processed in each iteration do not have fixed positions, and instead depend on when a collision happened in the previous iteration. We begin with a technical lemma, which proof is in the full version of this paper [11].

LEMMA 5.3. *Consider a random permutation of $1, 2, \dots, N$. Let M be the number of local minima in this permutation, i.e. the number of elements S_i smaller than both their neighbours S_{i-1} and S_{i+1} . Then, it holds that $\Pr[M \leq \frac{N-3}{18}] \leq \frac{24}{N-3}$.*

The next lemma will be the core of the lower bound proof. Similarly to Theorem 5.1, it analyses the collision probability for a cycle graph. Intuitively, it shows that for a prefix of edges of size $\Omega(\sqrt{N \cdot \log N})$ the probability of collision is almost 1. The (quite complex) proof is also provided in the full version of this paper [11].

LEMMA 5.4. *Consider a sequence of randomly shuffled cycle graph edges of size N and its prefix of size $W \geq C \cdot \sqrt{N \cdot \log N}$, where C is some constant which will be made clear in the proof. Suppose there are*

$M = pN$ pairs of adjacent colliding edges in this graph, $0 < p \leq \frac{1}{3}$ is a constant. Then the probability that there is no collision in this prefix is $O\left(\frac{1}{N^2}\right)$.

We can now finally state and prove our main lower bound result:

THEOREM 5.5. *The number of iterations of any parallel algorithm processing conflict-free prefix of edges (in particular, of the algorithm in Listing 2) is $\Omega\left(\frac{\sqrt{|E|}}{\sqrt{\log |E|}}\right)$ with high probability for a cycle graph.*

PROOF. Again consider a cycle graph. Initially, when the Union-Find data structure is empty, all connected components are of size and rank 1. In this case, we assumed that linking by rank or by size use random priorities to break ties. As a result, the number of vertices adjacent to two colliding edges is exactly the number of local minima in the random priorities permutation. Consequently, by Lemma 5.3, the number of pairs of colliding edges is at least $p|E|$ for some constant $p > 0$ with probability of failure $O\left(\frac{1}{|E|}\right)$.

We will prove that $\Omega\left(\frac{\sqrt{|E|}}{\sqrt{\log |E|}}\right)$ iterations are needed to process the first $\frac{p}{3}|E|$ edges. When an edge is processed, at most 2 pairs of colliding edges (adjacent to the edge) can stop colliding, since their adjacent connected components change. Thus, after $\frac{p}{3}|E|$ edges there will still be at least $p|E| - 2\frac{p}{3}|E| = \frac{p}{3}|E|$ pairs of colliding edges (which are still adjacent to connected components of size 1).

Let C be the constant from Lemma 5.4 for the case when there are at least $\frac{p}{3}|E|$ pairs of colliding edges. Let A_i be the event that in the i -th iteration, the processed prefix is of size at least $W = C \cdot \sqrt{|E|}\sqrt{\log |E|} + 1$. Let $I = \frac{p|E|}{W}$. Observe that to process $\frac{p}{3}|E|$ edges in less than $I = \Omega\left(\frac{\sqrt{|E|}}{\sqrt{\log |E|}}\right)$ iterations, at least one of the events A_i for these iterations should happen.

Now if we show that $Pr\left[\bigcup_{i \leq I} A_i\right] = O\left(\frac{1}{|E|}\right)$, the probability of using less than I iterations will be $O\left(\frac{1}{|E|}\right)$. For this we required about the existence of enough pairs of colliding edges, which is not satisfied with probability $O\left(\frac{1}{|E|}\right)$. Summing up these observations, with probability at least $1 - O\left(\frac{1}{|E|}\right)$ at least $\Omega\left(\frac{\sqrt{|E|}}{\sqrt{\log |E|}}\right)$ algorithm iterations are needed, completing the proof.

We first remark that

$$Pr\left[\bigcup_{i \leq I} A_i\right] \leq \sum_{i=1}^I Pr[A_i | \bar{A}_1, \dots, \bar{A}_{i-1}]$$

Consider the probability of such an event $Pr[A_i | \bar{A}_1, \dots, \bar{A}_{i-1}]$. Since the previous A_j ($j < i$) are assumed to not occur, less than $\frac{p}{3}|E|$ edges were processed before. The first edge in the next processing window may have had a collision with some edge in the previous iteration. However, no information is known about edges after it or their order, so we can still assume that the order is uniformly random. This lets us to apply Lemma 5.4 for the next $W - 1$ edges after the first one to obtain that the probability of no collision among these edges is $O\left(\frac{1}{|E|^2}\right)$, i.e., with probability at least $1 - O\left(\frac{1}{|E|^2}\right)$ no more than W edges can be processed in

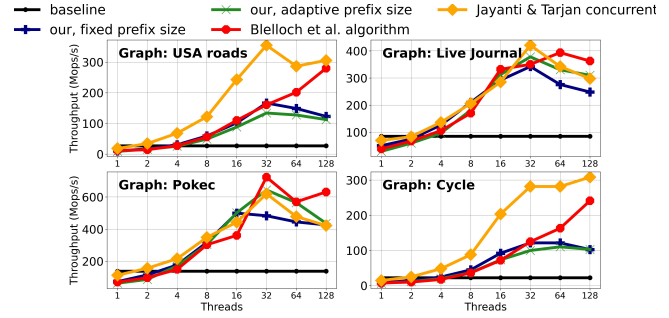


Figure 4: Performance comparison of parallel and concurrent Union-Find algorithms, relative to an optimized sequential baseline. Our algorithm is competitive with prior proposals with weaker guarantees.

i -th iteration due to data dependencies. Since it is true for any fixed possible start position of the processing window, it is also true in general, i.e., $Pr[A_i | \bar{A}_1, \dots, \bar{A}_{i-1}] = O\left(\frac{1}{|E|^2}\right)$. As a result, $Pr\left[\bigcup_{i \leq I} A_i\right] = O\left(\frac{I}{|E|^2}\right)$, which completes the proof once we substitute the value of I . \square

In other words, this theorem shows that, even when the algorithm in Listing 2 tries to process the whole edge sequence in each iteration, the number of iterations in the worst case is still asymptotically close to the work-efficient version from Theorem 4.7.

6 DISCUSSION

Overview. We have made new progress on understanding the parallel Union-Find problem, providing a new analysis suggesting that classic algorithms for this problem should parallelize well, since data races have low probability, which leads to the first algorithm that is both provably work-efficient and internally-deterministic.

Practical Performance. One interesting question is whether our algorithm is *practical*, i.e. whether ensuring internal determinism imposes significant performance bottlenecks. We investigate this question in detail in the full version of the paper [11]. Figure 4 briefly summarizes the results, showing that our algorithm can be competitive with that of Blelloch et al., which does not have theoretical guarantees, and with the fully-concurrent algorithm of Jayanti and Tarjan, which is not internally-deterministic.

Future Work. Our collision analysis can also be employed as a basis for analyzing the parallelization potential of Union-Find in other computation models, such as execution via relaxed scheduling [1]. Further, since our potential argument works over arbitrary graph families, our results could extend to the parallelization of other graph problems, such as minimum spanning tree (MST).

REFERENCES

[1] Dan Alistarh, Trevor Brown, Justin Kopinsky, and Giorgi Nadiradze. 2018. Relaxed Schedulers Can Efficiently Parallelize Iterative Algorithms. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing* (Egham, United Kingdom) (PODC '18). Association for Computing Machinery, New York, NY, USA, 377–386. <https://doi.org/10.1145/3212734.3212756>

- [2] Dan Alistarh, Alexander Fedorov, and Nikita Koval. 2019. In Search of the Fastest Concurrent Union-Find Algorithm. In *23rd International Conference on Principles of Distributed Systems, OPODIS 2019, December 17-19, 2019, Neuchâtel, Switzerland (LIPIcs, Vol. 153)*, Pascal Felber, Roy Friedman, Seth Gilbert, and Avery Miller (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 15:1–15:16. <https://doi.org/10.4230/LIPIcs.OPODIS.2019.15>
- [3] Daniel Anderson, Guy E Blelloch, and Kanat Tangwongsan. 2020. Work-efficient batch-incremental minimum spanning trees with applications to the sliding-window model. In *Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures*. 51–61.
- [4] Guy Blelloch, Jeremy Fineman, Phillip Gibbons, and Julian Shun. 2012. Internally Deterministic Parallel Algorithms Can Be Fast. *ACM SIGPLAN Notices* 47, 181–192. <https://doi.org/10.1145/2145816.2145840>
- [5] Guy Blelloch, Jeremy Fineman, and Julian Shun. 2012. Greedy Sequential Maximal Independent Set and Matching are Parallel on Average. *CoRR abs/1202.3205* (02 2012). <https://doi.org/10.1145/2312005.2312058>
- [6] Guy E. Blelloch. 2017. Some Sequential Algorithms Are Almost Always Parallel. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures* (Washington, DC, USA) (SPAA '17). Association for Computing Machinery, New York, NY, USA, 141. <https://doi.org/10.1145/3087556.3087602>
- [7] Guy E. Blelloch, Yan Gu, Julian Shun, and Yihan Sun. 2016. Parallelism in Randomized Incremental Algorithms. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures* (Pacific Grove, California, USA) (SPAA '16). Association for Computing Machinery, New York, NY, USA, 467–478. <https://doi.org/10.1145/2935764.2935766>
- [8] Laxman Dhulipala, Guy E. Blelloch, and Julian Shun. 2021. Theoretically Efficient Parallel Graph Algorithms Can Be Fast and Scalable. *ACM Trans. Parallel Comput.* 8, 1, Article 4 (apr 2021), 70 pages. <https://doi.org/10.1145/3434393>
- [9] Laxman Dhulipala, Changwan Hong, and Julian Shun. 2020. ConnectIt: A Framework for Static and Incremental Parallel Graph Connectivity Algorithms. *Proc. VLDB Endow.* 14, 4 (dec 2020), 653–667. <https://doi.org/10.14778/3436905.3436923>
- [10] Paul L. Erdos and Alfréd Rényi. 1984. On the evolution of random graphs. *Trans. Amer. Math. Soc.* 286 (1984), 257–257.
- [11] Alexander Fedorov, Diba Hashemi, Giorgi Nadiradze, and Dan Alistarh. 2023. Provably-Efficient and Internally-Deterministic Parallel Union-Find. arXiv:2304.09331 [cs.DS]
- [12] Manuela Fischer and Andreas Noever. 2019. Tight Analysis of Parallel Randomized Greedy MIS. *ACM Trans. Algorithms* 16, 1, Article 6 (dec 2019), 13 pages. <https://doi.org/10.1145/3326165>
- [13] Ashish Goel, Sanjeev Khanna, Daniel H. Larkin, and Robert Tarjan. 2014. Disjoint Set Union with Randomized Linking. *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms* (01 2014), 1005–1017. <https://doi.org/10.1137/1.9781611973402.75>
- [14] Siddhartha V. Jayanti and Robert E. Tarjan. 2021. Concurrent disjoint set union. *Distributed Comput.* 34, 6 (2021), 413–436. <https://doi.org/10.1007/s00446-020-00388-x>
- [15] Siddhartha V. Jayanti, Robert E. Tarjan, and Enric Boix-Adserà. 2019. Randomized Concurrent Set Union and Generalized Wake-Up. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, Peter Robinson and Faith Ellen (Eds.). ACM, 187–196. <https://doi.org/10.1145/3293611.3331593>
- [16] Anastasios Katsigiannis, Nikos Anastopoulos, Konstantinos Nikas, and Nectarios Koziris. 2012. An Approach to Parallelize Kruskal’s Algorithm Using Helper Threads. *Proceedings of the 2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops, IPDPSW 2012*, 1601–1610. <https://doi.org/10.1109/IPDPSW.2012.201>
- [17] Joseph B. Kruskal. 1956. On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem. *Proc. Amer. Math. Soc.* 7, 1 (1956), 48–50. <http://www.jstor.org/stable/2033241>
- [18] Anastasiia Postnikova, Nikita Koval, Giorgi Nadiradze, and Dan Alistarh. 2022. Multi-Queues Can Be State-of-the-Art Priority Schedulers. In *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Seoul, Republic of Korea) (PPoPP '22). Association for Computing Machinery, New York, NY, USA, 353–367. <https://doi.org/10.1145/3503221.3508432>
- [19] Julian Shun, Laxman Dhulipala, and Guy Blelloch. 2014. A Simple and Practical Linear-Work Parallel Algorithm for Connectivity. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures* (Prague, Czech Republic) (SPAA '14). Association for Computing Machinery, New York, NY, USA, 143–153. <https://doi.org/10.1145/2612669.2612692>
- [20] Julian Shun, Yan Gu, Guy E. Blelloch, Jeremy T. Fineman, and Phillip B. Gibbons. 2015. Sequential Random Permutation, List Contraction and Tree Contraction Are Highly Parallel. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms* (San Diego, California) (SODA '15). Society for Industrial and Applied Mathematics, USA, 431–448.
- [21] Natcha Simsiri, Kanat Tangwongsan, Srikanta Tirathapura, and Kun-Lung Wu. 2016. Work-Efficient Parallel Union-Find with Applications to Incremental Graph Connectivity. In *Euro-Par 2016: Parallel Processing*, Pierre-François Dutoit and Denis Trystram (Eds.). Springer International Publishing, Cham, 561–573.
- [22] Robert Endre Tarjan. 1975. Efficiency of a Good But Not Linear Set Union Algorithm. *J. ACM* 22, 2 (April 1975), 215–225. <https://doi.org/10.1145/321879.321884>