

In Search of the Fastest Concurrent Union-Find Algorithm

Dan Alistarh

IST Austria, Klosterneuburg, Austria
dan.alistarh@ist.ac.at

Alexander Fedorov

Higher School of Economics, St. Petersburg, Russia
JetBrains, St. Petersburg, Russia
afedorov2602@gmail.com

Nikita Koval

IST Austria, Klosterneuburg, Austria
JetBrains, St. Petersburg, Russia
ndkoyal@ya.ru

Abstract

Union-Find (or Disjoint-Set Union) is one of the fundamental problems in computer science; it has been well-studied from both theoretical and practical perspectives in the sequential case. Recently, there has been mounting interest in analyzing this problem in the concurrent scenario, and several asymptotically-efficient algorithms have been proposed. Yet, to date, there is very little known about the practical performance of concurrent Union-Find.

This work addresses this gap. We evaluate and analyze the performance of several concurrent Union-Find algorithms and optimization strategies across a wide range of platforms (Intel, AMD, and ARM) and workloads (social, random, and road networks, as well as integrations into more complex algorithms). We first observe that, due to the limited computational cost, the number of induced cache misses is the critical determining factor for the performance of existing algorithms. We introduce new techniques to reduce this cost by storing node priorities implicitly and by using plain reads and writes in a way that does not affect the correctness of the algorithms. Finally, we show that Union-Find implementations are an interesting application for Transactional Memory (TM): one of the fastest algorithm variants we discovered is a sequential one that uses coarse-grained locking with the lock elision optimization to reduce synchronization cost and increase scalability.

2012 ACM Subject Classification Theory of computation → Concurrent algorithms; Computing methodologies → Concurrent algorithms; Theory of computation → Graph algorithms analysis

Keywords and phrases union-find, concurrency, evaluation, benchmarks, hardware transactional memory

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2019.15

Related Version <https://arxiv.org/abs/1911.06347>

1 Introduction

Union-Find – the problem of maintaining the global connectivity structure of a set of integer elements based on their pair-wise connectivity – is a fundamental problem in computer science. This problem is alternatively known as Disjoint-Set Union (DSU), and is a key component of several classic algorithms, such as Kruskal’s [12] and Boruvka’s [6] algorithms for finding minimum spanning trees (MSTs), maintaining connected components and finding loops under edge additions, or finding strongly-connected components in directed graphs.

In its classic formulation, the sequential DSU problem assumes a finite ground set S , upon whose elements we perform the following operations:

- SameSet(u , v) checks whether two elements u and v are in the same set;
- Union(u , v) merges the sets to which u and v are currently assigned;



© Dan Alistarh, Alexander Fedorov, and Nikita Koval;
licensed under Creative Commons License CC-BY

23rd International Conference on Principles of Distributed Systems (OPODIS 2019).

Editors: Pascal Felber, Roy Friedman, Seth Gilbert, and Avery Miller; Article No. 15; pp. 15:1–15:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

- `Find(u)` returns the representative of the set in which `u` is located. This representative must be the same for all elements from the corresponding subset.

It is worth noting that in sequential case `SameSet` can be easily implemented via two `Find` invocations, on `u` and `v`, and checking whether the representatives coincide. Further optimizations exist to implement this method more efficiently; we discuss them in Section 4.2.

This classic problem is known to have a rich theoretical structure. A masterclass by Tarjan linked the upper bound on the worst-case time complexity of the problem to the inverse Ackermann function [18], followed by a matching lower bound for a restricted case [19], which was later extended to arbitrary algorithms to show optimality [13]. Later, Tarjan and Van Leeuwen [17] performed one of the first worst-case analyses for compaction heuristics in the context of Disjoint-Set Union (DSU), exposing the highly non-trivial fact that, with appropriate path compaction and linking heuristics, the problem can be solved in $\mathcal{O}(m \cdot \alpha(n, \frac{m}{n}))$ time complexity, where n is the number of elements, m is the number of operations, and α is a functional inverse of Ackermann’s function.

Patwary, Blair, and Manne [15] were the first to perform an in-depth experimental study of sequential strategies. Their work clearly showed the importance of minimizing the number of reads from memory on the performance of DSU, since the various strategies have very limited computational demands. Their experiments exposed the fact that the fastest sequential algorithm was the one designed by Rem in 1976 [8], which we describe in detail in the following sections.

The first concurrent algorithm for DSU was proposed by Cybenko, Allen, and Polito [7], whose key idea was using a spin-lock for write operations. Years later, Anderson and Woll proposed a *wait-free* concurrent algorithm [5], which is roughly a concurrent generalization of one of the linking-and-compression strategies studied by Tarjan and Van Leeuwen [17]. Their paper claims a worst-case upper bound for the algorithm of $\Theta(m \cdot (\alpha(m, 0) + p))$, where p is the number of parallel processes, and m is the number of operations; it was later observed that this proof is correct only under the non-standard assumption that threads cannot be preempted in between some certain operations in the algorithm [10].

Recently, Jayanti and Tarjan presented a set of correct and asymptotically-efficient concurrent DSU algorithms, which use fixed random priorities rather than ranks, based on the randomized sequential algorithm of [9]. These algorithms achieve a (total) work complexity upper bound of $\mathcal{O}(m \cdot (\alpha(n, \frac{m}{p}) + \log(\frac{np}{m} + 1)))$. Recent work by these authors, in collaboration with Boix-Adserà, suggested an algorithm with the same total work complexity bounds, but showed that this is optimal for a class of natural “symmetric algorithms” [11].

Our Contribution

Motivated by the significant recent interest in the concurrent DSU problem, as well as by its numerous practical applications, in this paper we perform the first thorough study of the practical performance of concurrent DSU implementations. While our focus is to implement and study existing algorithmic proposals, along the way we discover new optimizations and algorithmic insights. We start from the basic observation that, given the simple structure of most algorithms for DSU, memory access and synchronization costs will be the dominating factors behind practical performance. With this in mind, we analyze the performance of several classic baselines, and propose a host of optimizations to specifically reduce the impact of these factors.

We perform a range of experiments across several architectures, algorithm variants and optimizations, synchronization primitives, as well as workloads, to determine the fastest concurrent DSU algorithm. We provide a wide range of results and discuss our findings across several dimensions in detail in Section 5. In brief, the variant which appears to be “fastest” for most of the settings considered is an optimized sequential algorithm variant that

```

1 func ParallelMST(G: Graph) {
2   mst = {} // empty set initially
3   while mst.Size < THRESHOLD(G.Size) {
4     shortestEdges: Edge[G.Size] // shortest edges for each component
5     // Phase 1. Find the shortest edges for each vertex
6     parallel for e in G.Edges {
7       if SameSet(e.from, e.to) { G.Edges.remove(e); continue }
8       UpdateIfShorter(&shortestEdges[Find(e.from)], e)
9       UpdateIfShorter(&shortestEdges[Find(e.to)], e)
10    }
11    // Phase 2. Unite components and update MST
12    parallel for u in G.Vertices {
13      Unite(shortestEdges[u].from, shortestEdges[u].to)
14      mst += shortestEdges[u]
15    }
16  }
17  return mst + SequentialMST(graph)
18 }

```

■ **Listing 1** Parallel Boruvka’s algorithm pseudo-code used in our experiments.

leverages HTM (lock elision) for high path compaction while minimizing synchronization cost.

2 Experimental Setup

For our evaluation, we use two classic graph algorithms based on DSU data structures. The first maintains the *connected components* for a given graph. In this case, for benchmarking, we randomly split graph edges between threads, and also mark whether the Union of the SameSet operation should be executed with a given edge as a parameter. The set of graphs we consider is presented in Table 1.

The second benchmark is a parallel version of the *Boruvka’s algorithm* [6] for finding the minimum spanning tree. (The pseudo-code of the algorithm we use in the experiments is presented in Listing 1.) This algorithm performs at most $\log n$ iterations (where n is number of vertices), each is split into two phases. During the first phase, the algorithm finds the shortest adjacent edge for each vertex, removing the ones that connect vertices from the same component, the set of which is maintained by the DSU data structure. Here, the UpdateIfShorter function atomically checks whether the already stored edge is longer than the specified one, and replaces it in this affirmative case. During the second phase, the algorithm goes through all the representatives of the components and adds the corresponding shortest edges to the MST. The main idea of the concurrent algorithm is that both phases can be performed in parallel, with synchronization between them. However, if the number of remaining components is small, it is better to complete the work sequentially. Thus, we perform the parallel part until the already built MST part exceeds some size threshold. We consider only the parallel part in the benchmark.

Graph Inputs. We evaluate the algorithms above on a range of real-world and synthetic graphs. The connected components maintenance algorithm has linear time complexity (assuming for simplicity that the DSU operations work in constant time), while Boruvka’s algorithm requires logarithmic time to complete, which is why we test it on relatively smaller graphs. Nevertheless, both benchmarks use road and social network graphs, as well as random ones with similar properties. The list of graphs we use for the connected components and Boruvka’s algorithm benchmarks is presented in Tables 1 and 2 respectively.

The first graphs in both tables represent the USA road network, Central and West parts.

15:4 In Search of the Fastest Concurrent Union-Find Algorithm

■ **Table 1** Graphs for the connected components benchmark.

Graph	Vertices(M)	Edges(M)	Description
USA-ROADS	14.1	16.9	A graph with roads of the central part of the USA
LIVE-JOURNAL	4	34.7	A graph of ground-truth communities in LiveJournal social network
POKEC	1.6	30.5	A graph of friendship relations in Pokec social network
RANDOM	2.5	30	A random graph generated with the Erdős-Rényi model.
HIGH-CONTENTION	12	26.8	A synthetic graph for detecting “bad” DSU versions

■ **Table 2** Graphs for the Boruvka’s algorithm benchmark.

Graph	Vertices(M)	Edges(M)	Description
USA-ROADS	6.2	7.6	A graph with roads of the western part of the USA
BERKELEY-STANFORD	6.9	7.6	Hyperlinks between the Berkeley and the Stanford domains
INTERNET-TOPOLOGY	1.7	11.1	Internet topology graph, from traceroutes run daily in 2005
RANDOM	1	10	A random graph generated with the Erdős-Rényi model
HIGH-CONTENTION	4	8.9	A synthetic graph for detecting “bad” DSU versions

The next two graphs for the connected components benchmark are social network graphs, while the ones for the Boruvka’s benchmark are Internet topology graphs. Also, we use two large synthetic graphs. The first one is a randomly-generated Erdoes-Renyi graph, with a specified number of vertices and edges ($2.5M$ vertices and $30M$ edges for the connected components, and $1M$ vertices and $10M$ edges for the Boruvka’s algorithm); it is denoted as **RANDOM** in our experiments. The second synthetic graph is folklore for detecting “bad” DSU implementations in programming contests – most of the edges are incident to a small number of nodes; it is denoted as **HIGH-CONTENTION**.

Hardware. We perform the experiments on Intel, AMD, and ARM platforms; all the machines we used have several sockets, which induces Non-Uniform Memory Access (NUMA) effects. The detailed specifications are as follows:

- **Intel.** We used our local machine with 4 sockets, Intel Xeon Gold 6150 with 18 cores per socket, and hyperthreading enabled, for 144 hardware threads in total.
- **AMD.** We used a general-purpose Amazon AWS [1] instance with 6 sockets, 8-core AMD EPYC 7571 processors with hyperthreading enabled in each; 96 cores total.
- **ARM.** We used an instance on the Packet [3] cloud service with 2 sockets of 48-core ARM Cavium ThunderX processors; 96 cores total. The ARM memory model is more relaxed than the TSO one on Intel and AMD architectures, which would lead us to expect higher scalability.

Software. All algorithms and benchmarks are implemented either in Java or Kotlin, and compiled to the JVM byte-code; we use OpenJDK 11.0.4 with Ubuntu OS on all platforms. To avoid problems related to JIT compilers and reproducibility/benchmarking, we use the Java Microbenchmark Harness (JMH) library to run our benchmarks and collect statistics [2].

3 Sequential Implementations via Compressed Trees

We begin with an overview of sequential implementations, which are loosely based on the idea of maintaining a *compressed forest of trees* [18]. Each tree in this data structure corresponds to the membership of one set, where the root of the tree acts as its representative. The trees are implemented by maintaining parents for each element; thus, the data structure stores an array of parent links. Roots of the trees have their parent links point to themselves. Thus, in order to implement the `Find` operation, the algorithm “climbs” using parent links until it reaches the root. At the same time, the `Union` operation takes the corresponding tree roots as representatives, checks whether they coincide (finishing the operation in that case), and unites the sets by pointing one root to another. While the algorithm seems straightforward, it requires some heuristics to guarantee good time complexity.

Linking Strategies. When the `Union` operation decides to merge two different sets, it either changes the parent link of the first element to the second one or vice versa. Intuitively, we want to maintain the tree height as small as possible, so that `Find` operation works efficiently. The standard procedure is defining priorities on roots, so that the root with the lower priority is “hung” under the root with higher priority. The following definitions of priority are usually employed:

- *Tree size.* Each root maintains the size of its tree, and smaller roots are pointed to larger ones. This serves as a way of balancing the tree.
- *Rank.* Since we aim at optimizing the height, it is reasonable to store the height as a priority. When the ranks of the trees to be united are different, then the smaller tree is pointed to the larger one, and the ranks remain the same. However, when the ranks coincide, an arbitrary one is chosen as a new root, and its rank is incremented.
- *Random.* Another method to choose a new root is using a set of fixed random initial priorities; Tarjan et. al. first analyzed this technique in the sequential case [9], and after it Jayanti and Tarjan made an analysis for the concurrent case [10].

Path Compaction. Another way to make the data structure faster is making the tree “flatter” by shortening paths between nodes and the roots. In particular, notice that during the `Find` operation, the parent links can be changed to the higher ones without breaking correctness; this way, the algorithm reaches the root faster on the next `Find` invocation. Here are several strategies which are usually employed to compact trees:

- *Compression.* Once the root is found, all the elements on the search path can update their parent pointers to the root. The simplest implementation uses recursion and performs these updates in the last-in-first-out order. The disadvantage of this technique is that it requires performing two passes, from u to the root and back to u ; thus, in practice, this strategy produces extra cache misses.
- *Splitting.* This technique updates the parent links during the only traversal to the root. On each step, it reads both the parent and the grandparent of the current element, and updates its parent link to the grandparent. Thus, it compresses the paths from all visited elements to the root by one.
- *Halving.* This technique is similar to splitting, but after each step, each node is linked to the grandparent. Thus, parents of only half of the elements on the path are updated.

These compaction optimizations can guarantee $\mathcal{O}(\log n)$ amortized time complexity of the `Find` operation, while the linking techniques guarantee $\mathcal{O}(\log n)$ worst-case time complexity (probabilistic in case of random priorities). However, using both techniques at the same time, linking by priority and path compaction, guarantees $\mathcal{O}(m \cdot \alpha(n, \frac{m}{n}))$ time complexity on average, where n is the number of elements, m is the number of operations, and α is a functional inverse of Ackermann’s function.

15:6 In Search of the Fastest Concurrent Union-Find Algorithm

```
1 func Union(u, v) {
2   u = Find(u); v = Find(v)
3   if u == v: return
4   if rank[u] <= rank[v]:
5     parent[u] = v
6   else if rank[u] > rank[v]:
7     parent[v] = u
8   if rank[u] == rank[v]: rank[v]++
9 }
10 func Find(u) = while (true) {
11   p = parent[u]
12   gp = parent[p]
13   if p == gp: return p
14   parent[u] = gp
15   u = gp
16 }
```

■ **Listing 2** Sequential DSU with rank-based priorities using path compaction via halving.

Implementation Example. Listing 2 provides a pseudo-code of the DSU algorithm with ranks as priorities and path compaction via halving. Here we assume that the `SameSet` function is implemented with two `Find` invocations.

Evaluation. We evaluated all the combinations of linking and path compaction strategies. Similarly to the work by Patwary et al., we store priorities only for roots [15]. Thus, it is possible to use a single register to store either parent or priority – we only have to reserve one bit as a mark whether the element is a root or not; we show in the extended version of this paper [4] that this optimization significantly improves the performance.

In addition to the standard linking strategies, we suggest using a pseudo-random one, which essentially shuffles priorities in the range $1..n$ among n elements – the following well-known hashing formula due to Knuth is used:

$$\text{priority}(x) = (x + \text{SHIFT}) \times \text{BIG_PRIME} \bmod N.$$

It is worth noting that `SHIFT` can be chosen randomly to increase the algorithm’s robustness against some adversary to make the algorithm work slow.

Figure 1 shows the results of all the 12 combinations on both connected components and Boruvka’s algorithm benchmarks; we used our local Intel machine for this experiment.


Rank priorities are approximations for tree sizes. We predictably do not see significant difference between these linking strategies; however, using ranks requires updates only if two trees in the `Unite` operation have the same ranks. Thus, we do not consider the size-based strategy in further concurrent tests – we believe that it cannot outperform the rank one due to more memory updates.

The suggested pseudo-random priorities slightly outperform the real random ones in almost all scenarios (time of random generation was not considered). We believe that this is a consequence of the fact that we do not need to store pseudo-random priorities at all and that the pseudo-random method shuffles elements as evenly as possible. Thus, we use only pseudo-random priorities in further experiments. As for the comparison between pseudo-random and rank-based priorities, we do not see one clearly and consistently outperform the other in many scenarios. Therefore, we consider it important to test both these strategies in a concurrent environment. Similarly, we keep all linking strategies for further experimentation.

4 Evaluating Concurrent Implementations

4.1 Basic Variants

In accordance with the evaluation of the sequential algorithms, we consider only rank-based and pseudo-random linking strategies combined with all three path compression techniques. Similarly to the existing concurrent DSU implementations, we achieve atomicity by merging trees via `Compare-And-Set` (CAS) primitive.



■ **Figure 1** Time comparison for sequential DSU versions with different linking and path compaction strategies on the Intel machine on various input graphs. Lower is better.

Listing 3 contains a pseudo-code for the concurrent version with ranks based on the approach of Anderson and Woll [5]. For strategies with ranks we need to be sure that for a node a rank and a parent can not be updated concurrently, otherwise we can get a non-linearizable behavior, where, for instance, two nodes have parent links pointing to each other violating the forest structure invariant. We can create a structure storing both a parent and a rank and change with *Compare-And-Set* primitive the pointer to a structure, which was proposed by Anderson and Woll, but a faster way would be to store both a rank and a parent in the same register by either dividing it into two parts (i.e. a rank and a parent will have less bits) or by using the same trick as we used in the sequential case, when we stored priorities only for roots. The last technique was chosen since it uses twice less memory. A heuristic that was used for an optimization in the sequential case becomes important for the correctness in the concurrent case.

Figures 2 and 3 present the running time of the basic concurrent implementations for several of the strategies. One may expect that, following the sequential analysis, either rank-based or pseudo-random priorities would behave significantly better, but this is not apparent in the results, what motivates our further investigation. Another natural conjecture, which is apparent empirically, is that splitting and halving compaction strategies have very similar performance (since their cost is similar), while compression is inferior by comparison, due to two path traversals – for finding the root and for path compaction. Thus, we use only splitting heuristic in our next experiments. We also have collected statistics with the numbers of cache misses, which shows that the algorithms with the compression technique get about $\times 7$ more cache misses on loads and almost the same number on stores.

15:8 In Search of the Fastest Concurrent Union-Find Algorithm

```
1 func Union(u, v) = while (true) {
2   (u, ru) = Find(u)
3   (v, rv) = Find(v)
4   if u == v: return
5   if ru < rv {
6     if CAS(&A[u], u, v): return
7   } else if ru > rv {
8     if CAS(&A[v], v, u): return
9   } else { // ru == rv
10    if u < v && CAS(&A[u], u, v):
11      CAS(&A[v], rv, rv+1); return
12    if u > v && CAS(&A[v], v, u):
13      CAS(&A[u], ru, ru+1); return
14  }
15 }

16 func Find(u): (root, rank) {
17   p = A[u]
18   if isRank(p): return (u, p)
19   (root, rank) = Find(p)
20   if p != root:
21     CAS(&A[u], p, root)
22   return (root, rank)
23 }
24
25 func SameSet(u, v) = while (true) {
26   (u, _) = Find(u); (v, _) = Find(v)
27   if u == v: return true
28   if isRank(A[u]): // still a root?
29     return false
30 }
```

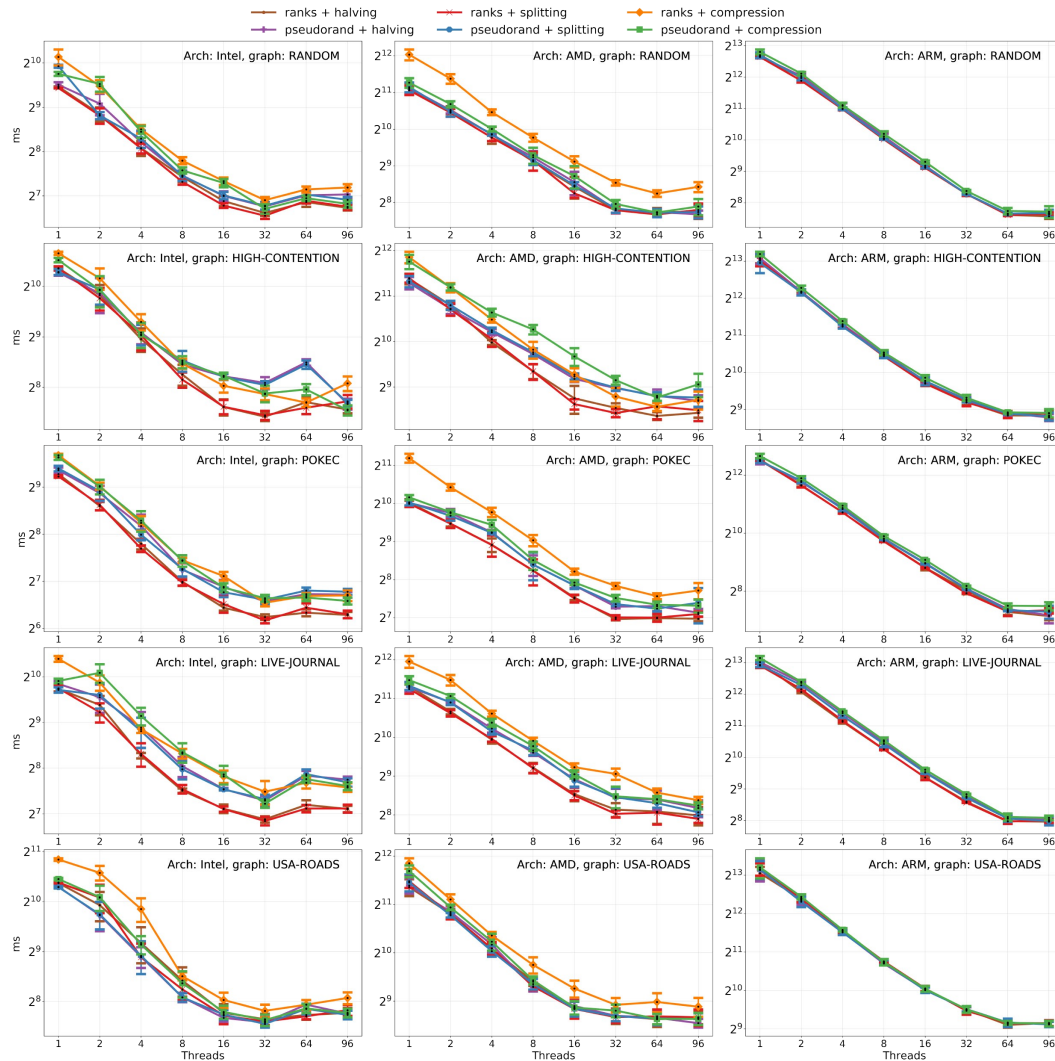
■ **Listing 3** Concurrent DSU with priorities via ranks and the path compression heuristic.

4.2 Optimizations

Path Compaction via Plain Writes. Figure 4 shows the average number of failed CAS operations for different algorithms and thread numbers on the connected component benchmark; we used our local Intel machine to collect these statistics. In worst-case scenarios, only 0.002% of the total number of CAS invocations fail. Thus, it should be safe to consider that there is almost no contention in practice, and there is no reason to perform several attempts to update parents, which differs from the theoretical worst-case [10].

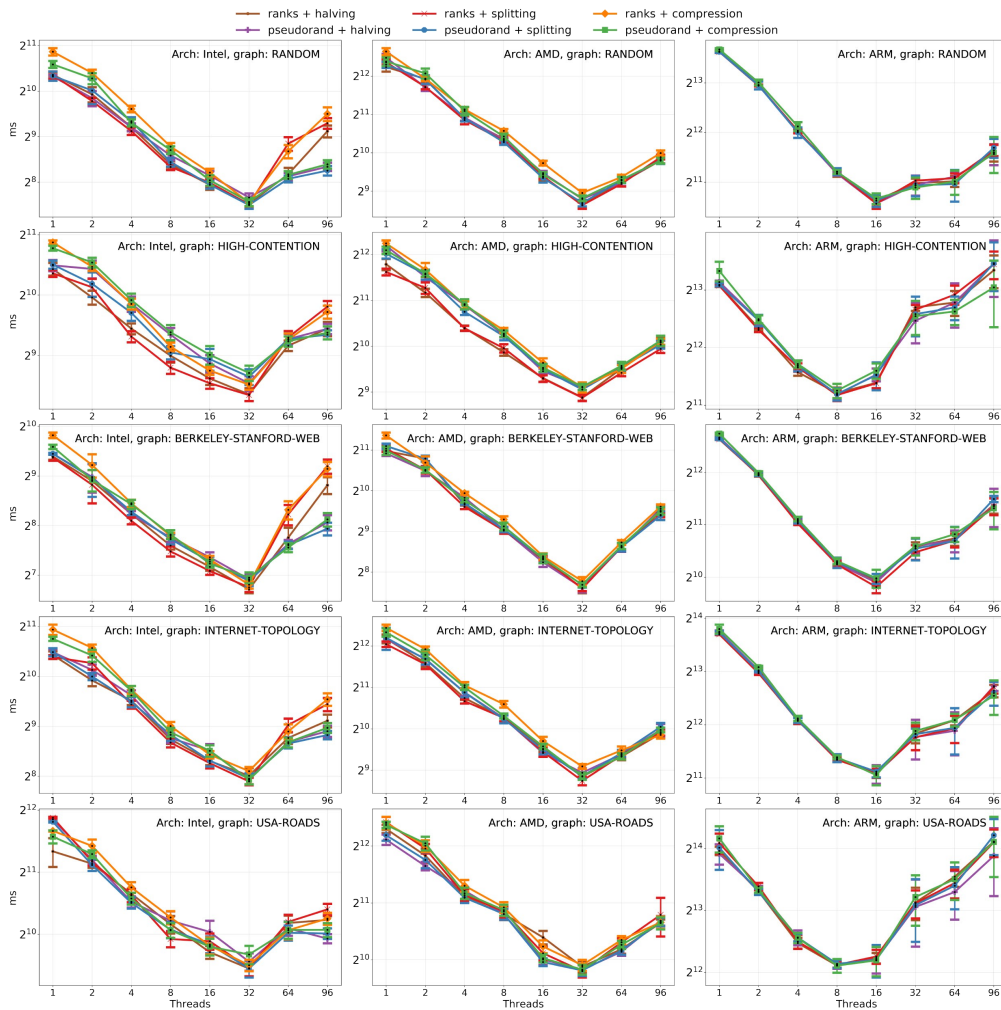
Since path compaction is a heuristic that influences performance but not correctness, we suggest trying to use simple writes (with or without memory barriers) for updating parent links. What is more, when using writes without memory barriers, we can also use reads without memory barriers at Find operations. While this trick is widely-known, our observation is that it does not break the correctness of the DSU algorithm. The only place where we need the reads with memory barriers is the moment when we check that a node is a root. If the check with the parent obtained from a plain read succeeded, then we should re-check using memory barriers. Figure 5 shows the comparison of versions with volatile and plain writes instead of CAS. While the version with volatile memory access does not show significant improvement, the one without memory barriers is faster to up to 40%. Thus, we consider it useful in almost all further experiments.

Early Recognition. The standard Union and SameSet implementations perform two Find operations on their parameters as the first step. However, if both elements are located in the same set, it could be more efficient to terminate the global operation when both Find invocations reach the lowest common ancestor; moreover, it potentially reduces the number of cache misses, which is especially important for NUMA architectures. Similarly, if these two elements are in different sets, it can also be detected during simultaneous climbs – it is guaranteed when the first Find reaches the root while the second one stays at an element with greater priority or vice versa. In this situation, SameSet can safely return false, while Unite can link the first root to the current element of the second Find climb. Goel et al. used this technique in sequential versions [9], while Jayanti and Tarjan adopted it for the concurrent environment [11]. Since we should know priorities for all nodes, we can not store them only for roots as we have done before for rank-based strategies; thus, we utilize twice more memory. Figure 6 shows that this optimization makes the algorithm faster for some scenarios, while the performance becomes worse or the same on others.




■ **Figure 2** Basic concurrent implementations comparison on the connected components benchmark. Each algorithm has been evaluated on five different graphs (see Table 1) and on Intel, AMD, and ARM platforms. Halving and splitting techniques show the best results, but the compression one is worse on both rank-based and pseudo-random strategies. While the trends on different architectures are similar, there are changes when NUMA effects appear, which can be noticed when the number of threads exceeds the number that one socket can have, thus forcing the algorithm to use several sockets. The versions on Intel and ARM machines do not scale after the moment they need to use two or more NUMA sockets, while the versions on AMD stop scaling only on three sockets. Another observation is that for ARM machines the difference between algorithms is less noticeable, because its processor is slower in comparison to other platforms, while the memory is as fast as in other machines.

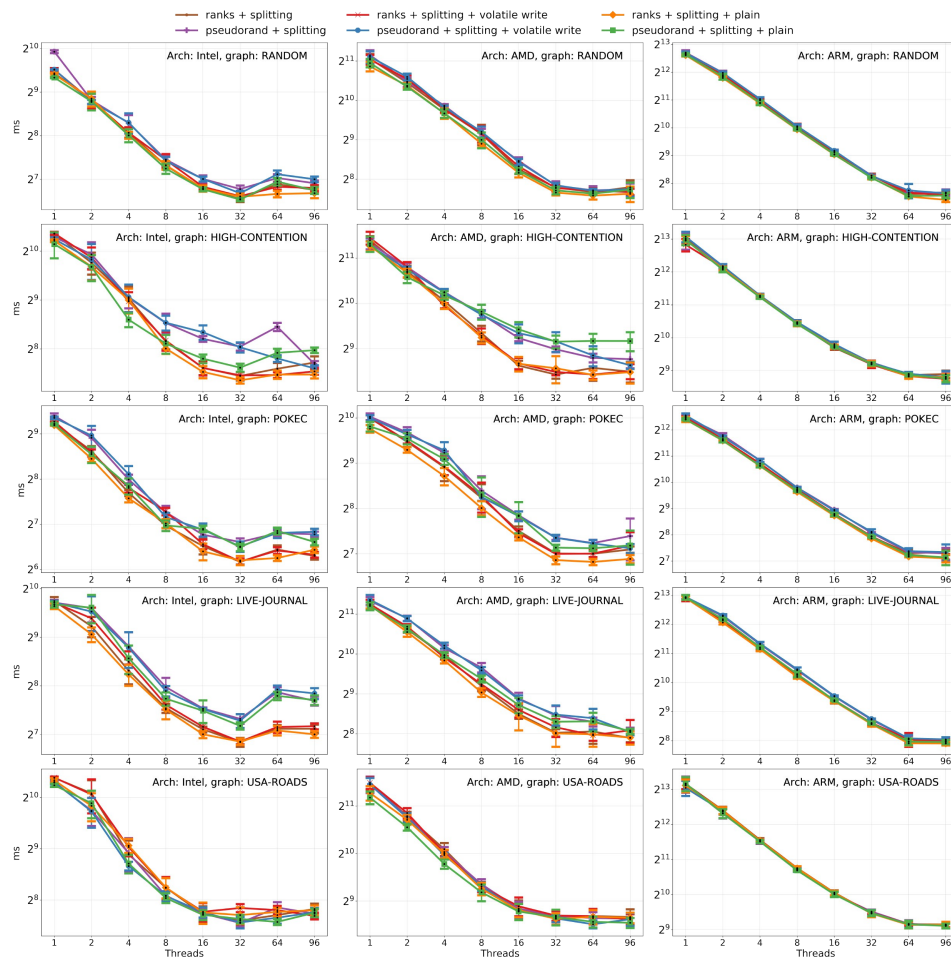
15:10 In Search of the Fastest Concurrent Union-Find Algorithm



■ **Figure 3** Basic concurrent implementations comparison on the Boruvka's algorithm benchmark. Here the algorithms do not scale well due to internal synchronization; thus, we are mainly interested in the fastest points – on 32 threads for Intel and AMD, on 8 or 16 threads for ARM.



■ **Figure 4** The average number of failed CAS operations on the connected components benchmark (Intel platform). Here we see that almost all CAS operations succeed; thus, in practice, there is no reason to perform several attempts to update parents.



■ **Figure 5** Volatile and plain writes instead of CAS operations to update parents. While volatile writes does not perform better than CAS operations, plain operations give a speedup up to 40%.

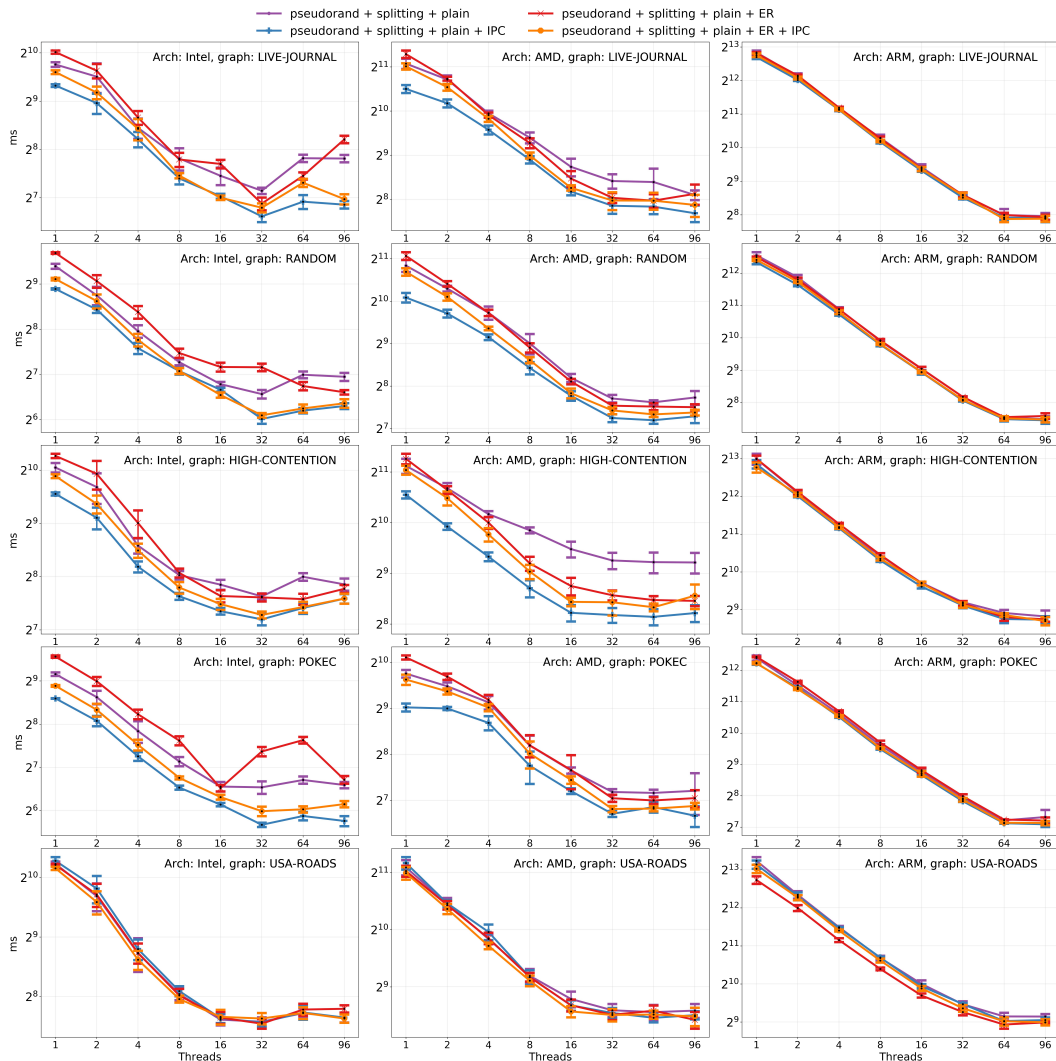
```

1 func Unite(u, v) = while (true) {
2   up = parent[u]; vp = parent[v]
3   if u == v or up == vp: return
4   if vp < up:
5     swap(&u, &v); swap(&vp, &up)
6   if u == up:
7     if CAS(&u.parent, u, vp):
8       return true
9   v = parent[up]
10  if up != v:
11    CAS(&u.parent, up, v)
12  u = up
13 }
14 func SameSet(u, v) = while(true) {
15   up = parent[u]
16   vp = parent[v]
17   if u == v or up == vp:
18     return true
19   if vp < up:
20     swap(&u, &v); swap(&up, &vp)
21   if u == up: return false
22   v = parent[up]
23   if up != v:
24     CAS(&u.parent, up, v)
25   u = up
26 }

```

■ **Listing 4** Concurrent Rem's disjoint-set union algorithm with path splitting.

Immediate Parent Check. Osipov et al. showed that it is likely for two elements to be united having the same parent links [14] (e.g., if the path is almost fully compressed after lots of Find invocations). Therefore, they suggested using an *immediate parent check (IPC)* optimization, which examines whether $\text{parent}[u]$ and $\text{parent}[v]$ are equal in the beginning of SameSet and Union operations and if they are, immediately returns. Figure 6 shows that *immediate parent check* significantly improves the performance of all the considered algorithms, especially on the Intel platform. However, the combination of both early recognition and immediate parent check optimizations work worse due to increasing the code complexity on some graphs and architectures. The *IPC* optimization is so efficient on our graphs due to having not many connectivity components; the best version presented in the figure is the one with it but without the *early recognition* optimization.



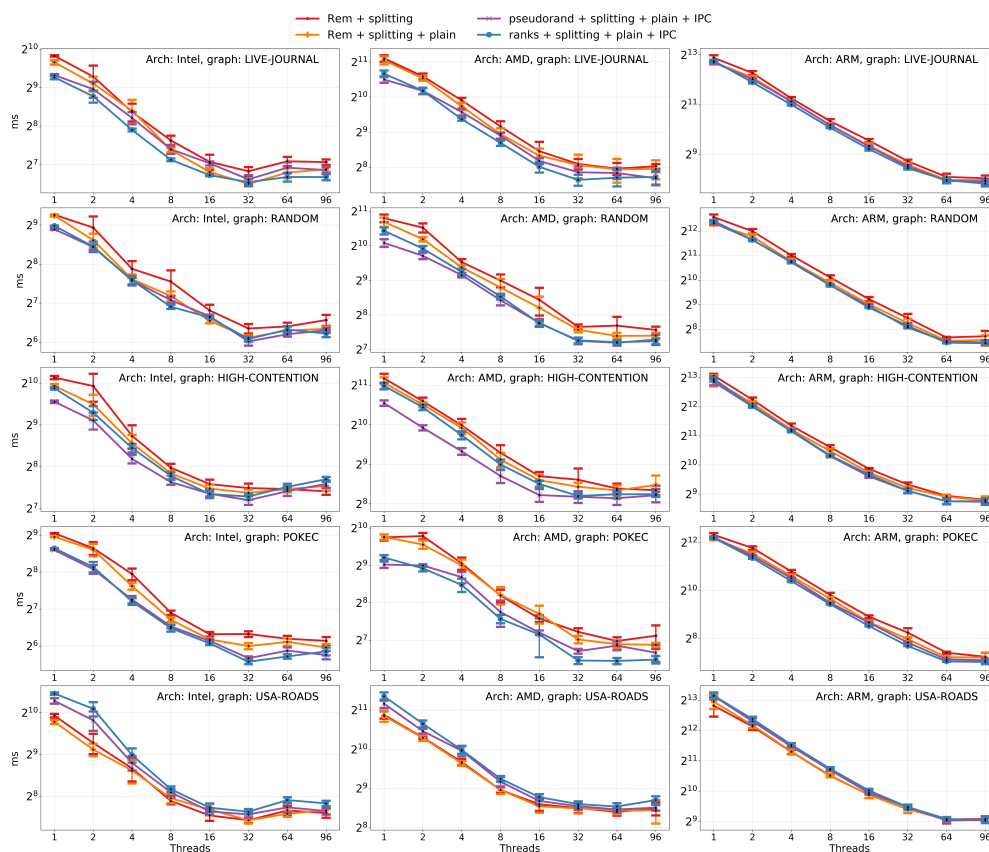
■ **Figure 6** Immediate parent check (IPC) and early recognition (ER) optimizations applied to algorithms with the pseudo-random linking strategy and path compaction via splitting on the connected components benchmark. Here we see that the variant with IPC and without ER wins in all scenarios. We omit the rank-based linking strategy and the Boruvka’s algorithm benchmark since the results are similar.

4.3 Rem's algorithm

Rem's algorithm was published in Dijkstra's book "A discipline of programming" [8]. To our knowledge, it is the first known interleaved algorithm, mixing early recognition and a specific linking strategy. The parents are the priorities of nodes. Then we do not need any memory optimizations, because we need nothing but parents. Similarly as in the early recognition algorithm, we perform two searches of roots, concurrently lifting at each step a node with a lower priority. A concurrent version of Rem's algorithm designed by us is presented at listing 4. Rem's algorithm was identified to perform best by Manne and Patwary in the sequential case [15]. While it is the best in the sequential environment, Figure 7 shows that the previously considered algorithms slightly outperform it.

Rem's algorithms are still better than most of the previously examined algorithms, and there are several reasons for it:

- They are based on the *early recognition* optimization.
- Since we need to compare priorities, we already know the parents at each step and can check them for equality having an automatic *immediate parent check* at each step.
- Their priorities are highly correlated with the distance to a root, which is important for *early recognition*. Random priorities never change, ranks change only during Union operations, while in Rem's algorithms priorities change after each path compaction, when a distance to the root decreases.

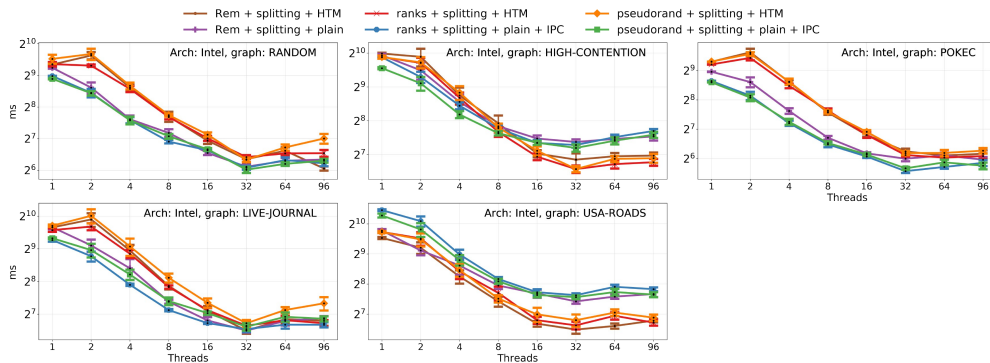


■ **Figure 7** Different variants of Rem's algorithm on the connected components benchmark compared with the previous best algorithms. While Rem's algorithm is the best in the sequential environment, the best of the previously examined ones are slightly better than its concurrent variants.

4.4 Transactional Memory

Since the number of failed CAS operations is imperceptible (see Figure 4), it is reasonable to expect that concurrent operations work with different memory locations almost all the time. Therefore, the union-find problem is potentially a good application for hardware transactional memory (HTM); we expect that almost all transactions will succeed. Instead of trying to use HTM as a fast-path for the concurrent versions, we check whether the best sequential algorithms can be scalable using a coarse-grained lock with the lock elision via HTM optimization [16]. The TM-based algorithms are coupled with all the sequential algorithms we discussed, and with all the corresponding sequential optimizations.

Figure 8 show the connected components benchmark results evaluated on our local Intel server; unfortunately, AMD and ARM do not support transactional memory at this point. The versions with HTM are definitely superior in some of the considered scenarios, being much simpler at the same time. The reason that for USA-ROADS and for HIGH-CONTENTION graphs on the connected components benchmark they perform better is that these graphs are sparse and operations are unlikely to touch the same memory. Therefore, we conjecture that as general rule algorithms with HTM perform better for sparse graphs. In the extended version of this paper [4] we performed an experiment on the Boruvka benchmark as well; due to more contention versions with plain writes and IPC optimization slightly outperform the HTM-based ones.



■ **Figure 8** Comparison of lock-based implementations with lock elision via HTM and the best previous algorithms on the connected components benchmark. The versions with HTM are significantly better on two of the considered scenarios and not much worse on other ones, being much simpler at the same time. Graphs where the results are better are sparse, so transactions are less likely to touch the same memory.

5 Discussion

We were the first to perform a thorough study of the practical performance of concurrent DSU implementations, exploring architectures (Intel, AMD, ARM), implementation variants (e.g., baseline, Rem), compaction strategies, as well as optimizations and synchronization techniques (lock-free, HTM). We mention the following salient points.

Memory Transfer Cost. One basic (and predictable) conclusion of our study is that one of the critical factors for the performance of DSU implementations is the amount of memory totally used and accessed per operation, and, in particular, the average number of cache misses per operation. This is natural since the computational cost of this workload is fairly negligible relative to the cost of the memory traversals. Thus, across architectures and

algorithmic variants, implementations that try to minimize memory accesses provide better performance. For instance, this is the case of the pseudo-random priorities optimization, which trades off a trivial amount of extra computation for reduced memory transfer.

Another point is that compression by rank (subtree height) appears to perform particularly well with respect to other methods (in particular, pseudo-random ranks) since it tends to produce higher-quality compressed trees. Notice that this diverges from the theoretical proposals of [10, 11], which avoid ranks due to additional algorithmic complexity.

Synchronization Cost. The second observation regards the synchronization cost of this problem: since most of the overhead of individual operations comes from synchronization cost (e.g., CAS for path compression), it is natural to investigate *relaxing* the strength of the synchronization primitives to improve performance. This is the primary motivation behind replacing CAS with plain or volatile write as an optimization, which does appear to lead to noticeable performance improvements. We note that, in this case, our empirical observations diverge from the theory; for instance, one of the algorithmic proposals of [10] requires a repeated CAS sequence to prove the performance upper bound, whereas our results clearly show that the simple update is sufficient in practice.

Compaction Quality. One of our objectives has been to investigate how the various heuristics for path compaction differ in terms of their practical performance. On this point, we find ourselves unable to adjudicate a clear winner between halving and splitting, which tend to provide similar performance. We do consistently observe that classic path compression leads to worse performance, which is to be expected due to the higher cost.

Scalability. In terms of scalability, we note that many graph algorithms are scalable up to the point where NUMA effects come into play, that is, at the point where threads might be accessing memory outside of their socket. This is especially salient in the case of the Boruvka application, where the cost of merging components across sockets dominates in the multi-socket case. We leave as an exciting future work direction the question of whether provably efficient concurrent NUMA-aware algorithms exist.

Performance Across Architectures. We notice that performance trends are consistent across architectures. Two noticeable differences are the better multi-socket performance of AMD processors (where NUMA effects become detrimental after 3 sockets are used as opposed to just 2 for Intel and ARM), and the fact that the plain write optimizations work better on Intel and AMD. The last observation befalls because ARM processors are slower, while the memory is as fast as in other machines, and we primarily optimize memory accesses.

Hardware Transactional Memory. The outstanding performance of HTM solutions was somewhat surprising for us. We explain this in hindsight by the simple fact that HTM allows the algorithm designer in this case to implement the simplest, most fast-path-efficient variant of the algorithm (in particular, Rem), which avoids most of the average synchronization costs of other techniques and allows more compiler optimizations, while contention is fairly limited in most scenarios.

Performance versus Monetary Cost. Since we ran our benchmarks on cloud instances, we are able to provide some intuition also with respect to the average cost per experimental run. In this case, the clear winner is the ARM machine, which offers significantly cheaper CPU time, at a very high core count.

The Fastest Concurrent Union-Find Algorithm. We conclude that, within the confines of our experimental setup, our optimized concurrent variants with either splitting or halving path compaction strategy and pseudo-random or rank priorities are best across architectures and inputs with an insignificant difference. This finding is perhaps disappointing from the theoretical perspective since we use plain writes instead of several attempts to update parents. Considering that we also use the immediate parent check optimization, it would be interesting to further study the guarantees of the resulting algorithms in the context of concurrent executions, perhaps for common distributions of (non-adversarial) input graphs.

References

- 1 Amazon Web Services (AWS) - Cloud Computing Services. <https://aws.amazon.com/>, 2019.
- 2 JMH - Java Microbenchmark Harness. <https://openjdk.java.net/projects/code-tools/jmh/>, 2019.
- 3 Packet Bare Metal - Cloud & Edge Computing Infrastructure Provider. <https://www.packet.com/>, 2019.
- 4 Dan Alistarh, Alexander Fedorov, and Nikita Koval. In Search of the Fastest Concurrent Union-Find Algorithm. *arXiv preprint arXiv:1911.06347*, 2019.
- 5 Richard J. Anderson and Heather Woll. Wait-free Parallel Algorithms for the Union-Find Problem. In *In Proc. 23rd ACM Symposium on Theory of Computing*, pages 370–380, 1994.
- 6 Otakar Boruvka. O jistém problému minimálním, 1926.
- 7 George Cybenko, T. G. Allen, and J. E. Polito. Practical parallel Union-Find algorithms for transitive closure and clustering. *International Journal of Parallel Programming*, 17:403–423, 1988.
- 8 E.W. Dijkstra. *A discipline of programming*. Prentice-Hall series in automatic computation. Prentice-Hall, 1976.
- 9 Ashish Goel, Sanjeev Khanna, Daniel H. Larkin, and Robert Tarjan. Disjoint Set Union with Randomized Linking. *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1005–1017, January 2014. doi:10.1137/1.9781611973402.75.
- 10 Siddhartha Jayanti and Robert Tarjan. A Randomized Concurrent Algorithm for Disjoint Set Union. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016*, pages 75–82, July 2016. doi:10.1145/2933057.2933108.
- 11 Siddhartha V. Jayanti, Robert E. Tarjan, and Enric Boix-Adserà. Randomized Concurrent Set Union and Generalized Wake-Up. In *PODC*, 2019.
- 12 Joseph B Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society*, 7(1):48–50, 1956.
- 13 Michael L. Fredman and Michael Saks. The Cell Probe Complexity of Dynamic Data Structures. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*, pages 345–354, January 1989. doi:10.1145/73007.73040.
- 14 Vitaly Osipov, Peter Sanders, and Johannes Singler. The Filter-Kruskal Minimum Spanning Tree Algorithm. In *ALENEX*, 2009.
- 15 Md Mostofa Ali Patwary, Jean Blair, and Fredrik Manne. Experiments on union-find algorithms for the disjoint-set data structure. In *International Symposium on Experimental Algorithms*, pages 411–423. Springer, 2010.
- 16 Ravi Rajwar and James R Goodman. Speculative lock elision: Enabling highly concurrent multithreaded execution. In *Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture*, pages 294–305. IEEE Computer Society, 2001.
- 17 Robert E. Tarjan and Jan van Leeuwen. Worst-case Analysis of Set Union Algorithms. *J. ACM*, 31:245–281, 1984.
- 18 Robert Endre Tarjan. Efficiency of a Good But Not Linear Set Union Algorithm. *J. ACM*, 22(2):215–225, April 1975. doi:10.1145/321879.321884.
- 19 Robert Endre Tarjan. A class of algorithms which require nonlinear time to maintain disjoint sets. *Journal of computer and system sciences*, 18(2):110–127, 1979.