

---

# A SCALABLE CONCURRENT ALGORITHM FOR DYNAMIC CONNECTIVITY

---

**Alexander Fedorov**  
JetBrains Research and Higher School of Economics  
✉

**Nikita Koval**  
JetBrains Research  
✉

**Dan Alistarh**  
IST Austria  
✉

May 19, 2021

## ABSTRACT

*Dynamic Connectivity* is a fundamental algorithmic graph problem, motivated by a wide range of applications to social and communication networks and used as a building block in various other algorithms, such as the bi-connectivity and the dynamic minimal spanning tree problems. In brief, we wish to maintain the connected components of graph under dynamic edge insertions and deletions. In the sequential case, the problem has been well-studied from both theoretical and practical perspectives. However, much less is known about efficient *concurrent* solutions to this problem. This is the gap we address in this paper.

We start from one of the classic data structures used to solve this problem, the Euler Tour Tree. Our first contribution is a non-blocking *single-writer* implementation of it. We leverage this data structure to obtain the first truly concurrent generalization of dynamic connectivity, which preserves the time complexity of its sequential counterpart, but is also *scalable* in practice. To achieve this, we rely on three main techniques. The first is to ensure that *connectivity queries*, which usually dominate real-world workloads, are non-blocking. The second non-trivial technique expands the above idea by making all queries that do not change the connectivity structure *non-blocking*. The third ingredient is applying fine-grained locking for updating the connected components, which allows operations on disjoint components to occur in parallel.

We evaluate the resulting algorithm on various workloads, executing on both real and synthetic graphs. The results show the efficiency of each of the proposed optimizations; the most efficient variant improves the performance of a coarse-grained based implementation on realistic scenarios up to 6x on average and up to 30x when connectivity queries dominate.

## 1 Introduction

*Dynamic connectivity* is a fundamental algorithmic problem on graphs under edge updates, motivated by a wide range of applications in the context of social and communication networks. It is also used as a key building block for various other algorithms, such as the bi-connectivity and the dynamic minimal spanning tree problems. Essentially, the goal is to maintain a query on whether two vertices are in the same connected component under edge insertions and deletions. Thus, the following operations on undirected graph  $G$  should be supported:

- $E(u, v)$  adds a new edge  $(u, v)$  to  $G$ ;
- $R(u, v)$  removes the edge  $(u, v)$  from  $G$ ;
- $Q(u, v)$ :  $\text{bool}$  checks whether  $u$  and  $v$  are in the same component of connectivity, returning  $\text{true}$  in this case.

We refer to both  $\text{insert}(u, v)$  and  $\text{delete}(u, v)$  as *write operations* or *modifications*, and to  $\text{query}(u, v)$  as a *read operation* or *query*. Intuitively, read operations are easier to parallelize, while modifications require non-trivial synchronization.

One classic approach to solve the dynamic connectivity problem is to use *Euler Tour Trees (ETT)* [1] to maintain a spanning forest of the graph. This provides a relatively simple way to implement both edge additions and connectivity queries in  $\mathcal{O}(\log N)$  time per operation, where  $N$  is the number of graph vertices. However, edge removals are non-trivial. The challenge is that, when a spanning edge is removed, there can be an edge that reconnects the spanning trees; this edge needs to be identified and added to the spanning forest. The classic algorithms proposed by Henzinger et al. [1] and Holm et al. [2] partition non-spanning edges into levels, to achieve amortized polylogarithmic time complexity per modification. In turn, these algorithms were improved to  $\mathcal{O}(\log N (\log \log N)^2)$  and  $\mathcal{O}(\frac{\log^2 N}{\log \log N})$  time complexity in [3] and [4], respectively, using similar ideas. Interestingly, these algorithms are close to the theoretical lower bounds [5, 6].

Even though this problem is well-studied in the sequential world in both theoretical and practical aspects [7], to our knowledge there is still no truly concurrent algorithm with polylogarithmic time per operation. Nonetheless, we highlight a few approaches which take advantage of parallelism for variants of this problem.

Aksenov et al. introduced a *parallel combining* technique [8], which allows performing read operations in parallel. In contrast, all modifications are performed sequentially and cannot be executed in parallel with reads. To evaluate this technique, they applied it to the dynamic connectivity problem so that concurrent  $\text{query}(u, v)$  invocations can be executed in parallel — the idea is similar to the readers-writer lock. However, this is essentially a sequential algorithm which leverages the *flat combining* technique [9].

Acar et al. [10] proposed an algorithm which handles operations *in batches*, allowing some operations in a batch to be processed in parallel and therefore more efficiently. One shortcoming of this approach is that it requires requests to be batched, and, as such, it is not a truly concurrent solution.

**Our Contribution.** In this paper, we present the first concurrent solution for the dynamic connectivity problem. Our algorithm is *linearizable* [11]. Similarly to many sequential dynamic connectivity algorithms, we use *Euler Tour Trees* as the backbone structure, where each tree represents one connected component. Instead of making the ETT algorithm fully concurrent, we present a *single-writer* solution which implements the  $\text{insert}(u, v)$  operation in a non-blocking way; hence, all the modifications should be guarded by a lock. In short, we partition *ETT* merge and split operations into their logical and physical components, so these operations appear atomic to concurrent readers, regardless of internal cuts and links. To make concurrent reads linearizable, we introduce versioning for the roots of the Cartesian trees [12] — which represent Euler Tour Trees; please see Section 3 for a detailed overview. These versions help to detect connectivity changes when snapshotting roots for the connectivity check. The ideas we introduce can be employed for developing a concurrent single-writer version of any balanced tree which supports the standard  $\text{insert}(A, B)$ ,  $\text{delete}(T)$ ,  $\text{query}(u, v)$ ,  $\text{insert}(u, v)$ , and  $\text{query}(u, v)$  operations.

Using the algorithm of Holm et al. [2] for the dynamic connectivity problem and these single-writer *ETT*-s, we automatically get a single-writer dynamic connectivity algorithm. We find the “single-writer” restriction reasonable in this application, since it allows us to implement Euler Tour Trees almost as efficiently as in the sequential case, while allowing queries, which are usually the most frequent operations, to run in parallel. Furthermore, this assumption reduces the overall complexity compared to a potential fully lock-free algorithm, which is bound to involve extremely complex book-keeping of components. Specifically, imagine a fully lock-free generalization of the algorithm of Holm et al. for the dynamic connectivity problem, e.g., implemented via lock-free software transactions. Since a replacement search takes  $\Omega(N)$  in the worst case, and other operations can cause it to fail and re-run, the upper bound on the total work would be  $\Omega(MN)$ , which is substantially worse than  $\mathcal{O}(M \log^2 N)$  of the sequential solution, where  $M$  is the number of operations.

Next, to improve scalability, we discuss how fine-grained locking can be applied to our proposal, so that updates on different components can be naturally parallelized. Surprisingly, this variant also requires our single-writer *ETT*, even when all operations, including connectivity queries, are executed under locks. A further optimization we propose gives a way to perform edge additions and removals that do not change the spanning forest in a non-blocking way.

Finally, we evaluate the resulting algorithm on various workloads, both on real-world and synthetic graphs. The results show the efficiency of each of the proposed optimizations; the most efficient variant improves the performance of a coarse-grained locking implementation on realistic scenarios up to 6x on average and up to 30x when connectivity queries dominate the workload.

This work has been published at SPAA ’21 and is available as [13].

## 2 Related Work

**Sequential Algorithms.** Variants of the connectivity problem have a rich history in the sequential algorithm literature, e.g. [14, 15, 16]. The first polylogarithmic-time dynamic connectivity algorithm was proposed by Henzinger and King [1]. The main idea is to maintain a spanning forest using *Euler Tour Trees (ETTs)* and partition the rest of the graph into  $\mathcal{O}(\log N)$  levels. We discuss Euler Tour Trees in detail in Section 3. Whenever removing spanning edges, they employ a *random sampling* technique to find an edge that can reconnect the trees. As a result, updates achieve randomized  $\mathcal{O}(\log^3 N)$  work on average per operation, while  $\text{find}$  queries require only  $\mathcal{O}(\frac{\log N}{\log \log N})$  work.

A faster and arguably simpler *deterministic* algorithm was introduced by Holm, Lichtenberg, and Thorup [2]. It is based on the same ideas as the algorithm of Henzinger and King and uses nested levels of edges to get deterministic amortized  $\mathcal{O}(\log^2 N)$  per operation and the same  $\mathcal{O}(\frac{\log N}{\log \log N})$  time complexity for reads. Later, Thorup decreased the amount of used memory to  $\mathcal{O}(N + M)$  [17], where  $M$  is the number of operations.

Further algorithms proposed additional enhancements. First, in 2011, Wulff-Nilsen introduced the *shortcut structure* and *lazy local trees* to optimize modification time complexity by a  $\log \log N$  factor [3]. Second, in 2015, Huang et al. applied these ideas to the Henzinger and King algorithm, achieving  $\mathcal{O}(\log N (\log \log N)^2)$  amortized modification time and  $\mathcal{O}(\frac{\log N}{\log \log \log N})$  query time.

As for the algorithms optimizing worst-case time, the best known algorithm was proposed by Wang in 2015 [18]. It employs the *cutset structure* to get  $\mathcal{O}(\log^4 N)$  modification time with an exponentially small probability of an error. The existence of a deterministic algorithm with polylogarithmic worst-case time per operation is still open. The best such algorithm has  $\mathcal{O}(\frac{\sqrt{N}}{w^{1/4}})$  time by combining the *square root decomposition* and *bit compression*, where  $w$  is the machine word size.

Patrascu and Demaine proved that the lower bound on operation time complexity is  $\Omega(\frac{\log N}{\log \log n})$  [5]. Furthermore, Patrascu and Thorup showed that the modification time can not be  $o(\log N)$ , unless the query time is  $\Omega(N^{1-o(1)})$  [6]. Consequently, all presented polylogarithmic algorithms can be seen as “nearly-optimal.”

**Parallel Algorithms.** Acar et al. in 2019 suggested employing parallelism to process batches of operations of the same type [10]. Their algorithm has no total work overhead due to multiple threads and becomes asymptotically faster with larger batch sizes, however, even for batches of size  $n^{1-\varepsilon}$  there still will be no asymptotic difference, where  $\varepsilon > 0$  is any constant. Another problem is the requirement to have large batches of queries of the same type. Meanwhile, if there is a *concurrent* algorithm, it is trivial to use it to process batches in parallel, which means that a *concurrent* algorithm can solve more general problems.

To construct a concurrent algorithm, Aksenov et al. applied the *flat combining* technique [9] to the Holm et al. algorithm in [8]. However, this technique’s standard implementation does not provide much better performance than the coarse-grained locking algorithm. Therefore, they introduced its read-optimized version named *parallel combining*, which processes query operations in parallel and scales better than the standard flat-combining on read-dominated scenarios.

A concurrent lock-free algorithm was proposed by Chatterjee et al. [19]. They develop a BFS-based connectivity query with linear time complexity per operation, which is inferior to the polylogarithmic time of the sequential algorithms.


As far as Euler Tour Trees are concerned, Tseng et al. designed a batch-parallel ETT algorithm using a phase-concurrent skip-list [20]. *Phase-concurrent* means that no concurrent operations of different types are allowed. However, in our algorithm we focus on allowing concurrent reads and updates because reads are fast and frequent, while modifications are slow and rare in realistic read-dominated scenarios. Intuitively, in such scenarios, slow and rare operations should not cause fast and frequent operations to wait.

Compressed functional trees with merge and union operations were used by Dhulipala et al. for low-latency graph streaming [21]. They get concurrent reads “for free”, since they employ *immutable* data structures. Similarly to Tseng et al. [20] and Acar et al. [10], they use parallelism to process batches of operations. Nevertheless, in our problem trees should have links both from parents to children and from children to parents, because connectivity queries traverse trees from bottom to top. This makes batch-parallel processing and partially-persistent data structures much more difficult, so the same ideas can not be applied easily to the dynamic connectivity problem.

## 3 Single-Writer Euler Tour Tree

As a first step, we present a concurrent single-writer multiple-reader lock-free *Euler Tour Tree (ETT)* algorithm.

**Sequential Euler Tour Tree.** In short, *ETT* specifies an undirected rooted tree as a graph traversal that starts and ends in the root and covers all edges exactly twice — see the picture on the right as an example. This traversal is exactly the same as in the depth-first search (DFS) graph algorithm started from the root.



Henzinger and King solved the dynamic connectivity problem by representing each component via an *ETT* [1]. Consider the traversal path from the picture above, which stores all the edges and the first occurrence of each vertex:

$[v_1, 1, v_2, 2, v_3, 2, 3, v_4, 4, v_5, 4, 5, v_6, 5, 3, 1]$

Assume we want to split the original tree by removing the edge “3” and produce two *ETT*-s as a result. The edge “3” cuts the traversal path into three parts. Notice that the blue part in the center and the linked red parts on the ends produce the *ETT*-s we need. Thus, we can store these traversals in a data structure with  $\cup$  and  $\cap$  operations, so split is performed by two  $\cup$ -s and one  $\cap$ . For reasons described later, we use Cartesian trees [12] with implicit keys, where every node has a random priority, and the tree forms a heap according to these priorities. This data structure supports  $\cup$  and  $\cap$  operations in  $\mathcal{O}(\log N)$  expected time.


This way, *ETT*-s solve the dynamic connectivity problem for forests (sets of trees) in the following way:

- $\text{find}(u, v)$  finds the corresponding Cartesian tree roots starting from the nodes  $u$  and  $v$  and checks them for equality;
- $\text{merge}(u, v)$  merges the corresponding Cartesian trees;
- $\text{split}(u, v)$  splits the Cartesian tree in the way described above.

All these operations work in  $\mathcal{O}(\log N)$  time.

**Atomic Merge and Split.** We want to construct a linearizable *ETT* implementation that allows  $\text{merge}$  and  $\text{split}$  operations to be executed by only one thread and the  $\text{find}$  query to be executed by multiple threads in a non-blocking way. The idea is to use the same sequential algorithm, when  $\text{find}(u, v)$  finds the roots of  $u$  and  $v$  in the corresponding Cartesian trees. However, we need to do both searches atomically under concurrent updates.


The first problem on the way to such an implementation is the appearance of out-of-thin-air components during *ETT* splits and merges. Consider an edge addition between  $AB$  and  $CD$  components on Figure 1. To merge the corresponding *ETT*-s, the Cartesian trees should be cut by the edge endpoints and then linked in the presented way. However, when  $\text{find}(u, v)$  is in the middle of the update and the cuts are finished but the trees are not linked yet, concurrent readers may see non-existing connectivity components, appeared out of nowhere.



**Figure 1:** Edge addition between  $AB$  and  $CD$  components. The corresponding Cartesian trees should be cut by the edge endpoints, followed by linking the resulting four parts in a proper order. However, in the center state concurrent readers may see non-existing connectivity components.


In order to solve this problem, we partition *ETT* splits and merges into *logical* and *physical* parts. The logical part consists only of one instruction and will be the linearization point of the operation.

The edge addition (merge) procedure starts with linking one Cartesian tree to another by setting the parent link — this way; the trees seem already merged for concurrent readers. It is correct since for the Cartesian trees we know the node that will become the common root – the root that has higher priority. Note that parent links always lead to a node with higher priority, so the parent link graph is acyclic and there is only one sink node after the first step. The following cuts keep the parent links, so the tree stays merged during the restructuring. Figure 2 illustrates this algorithm.



**Figure 2:** *Atomic* edge addition between AB and CD components. At first, the corresponding trees are merged logically. The tree should then be restructured, but the parent links of the subtrees produced by cuts are kept.

The edge removal (split) procedure first prepares the tree to be split and then applies the operation by a single parent link change. It means that we at first cut the Cartesian tree into three parts, *A*, *B*, and *C*, but keep the parent link as we did for merging. After that, the first and the third trees, *A* and *C*, are united. The constructed *AC* and *B* trees correspond to the *ETT*-s we need to have at the end. Notice that either the parent of the root of *B* is a node in *AC* or vice versa. The corresponding parent link removal splits the *ETT*-s both physically and logically. Figure 3 illustrates the described procedure.



**Figure 3:** *Atomic* edge removal. The algorithm keeps the parent links during the required cuts and merges, and the last unlinking applies the operation both logically and physically.

**Linearizable connected(*u*, *v*) Operation.** Once the out-of-thin-air problem is solved, we can construct a linearizable algorithm for  $\text{connected}(u, v)$ . The idea is to maintain versions of the Cartesian tree roots and change them on every split and merge operation. We want to use these versions to check that roots for *u* and *v* were snapshotted correctly by re-checking these versions.

Yet, this approach immediately meets a problem — it is non-trivial to perform the logical part of the operation and change the root version atomically since these actions change different memory locations. For instance, when an edge addition changes the parent link that belongs to the root with lower priority, the version increment should be performed on the root with higher priority (see Figure 2 that illustrates the edge addition procedure).

As a solution, we suggest updating the root version *before* each merge and split procedure; thus, before each edge addition and removal. Since there is a single writer, it is guaranteed that the version is *at most* one step ahead for concurrent readers. However, this relaxation requires a non-trivial check to ensure that the obtained roots were snapshotted atomically.

If version and parent link updates were atomic, then the connectivity check would find the Cartesian tree roots for both *u* and *v*, and then find the root for *u* again. When both roots and the corresponding versions found for *u* coincide, it is guaranteed that the snapshot is correct, and the connectivity check should simply compare the obtained roots.

With the versions that can be one step ahead, we may need more re-checks. When roots for *u* and *v* coincide and the second root search for *u* returns the same node and version, there was no change between the first root search for *u* and the root search for *v* that affected the connectivity check.

However, when the obtained roots for *u* and *v* are different, and the second root search for *u* returns the same node and version, we cannot be sure that *u* and *v* are in different components. Despite this observation, an additional pair of checks that the roots of *u* and *v* and their versions has not been changed solves the problem. The corresponding proof is

discussed later in this section. In addition, we show that the algorithm is incorrect without the last root search and the corresponding re-check in Appendix A.

The described algorithm is presented in Listing 1.

```

1 func find_u_v: B0= while (true) {
2 // Get the roots of 'u' and 'v' with the versions
3 (u_root := find_u)
4 (v_root := find_v)
5 // Has the component of 'u' been changed?
6 if (u_root != find_u): continue
7 if u != v
8 // 'u' and 'v' are likely to be in different
9 // components, re-check that 'u_root' and
10 // 'v_root' were snapshotted atomically
11 if (u_root != find_v): continue
12 if (v_root != find_u): continue
13 return u == v
14 }
15 // Returns the root with the current version
16 func find_root: (B I) {
17 cur := u
18 while (true):
19     par := cur // read the parent of 'cur'
20     if par == null: break // is 'cur' the root?
21     cur = par // go towards the root
22 // Return the found root with the current version
23 return (par cur)
24 }

```

**Listing 1:** Linearizable connectivity check. `find_u_v` follows parent links until it finds the root and its version.

**Proof of Correctness.** To prove the correctness of the `find_u_v` algorithm we first prove the following lemma:

**Lemma 3.1.** *Assume that `find_root` returned  $(r, v_1)$  for some vertex and another `find_root` returned  $(r, v_2)$ . Consider the component of connectivity with vertex  $r$ . If a modification in it started after the first `find_root` completion and finished before the root is found in the second `find_root` invocation, then the root version changed, i.e.  $v_1 \neq v_2$ . By “root finding” we mean the moment when `find_root` reads `cur.parent` that equals `null`.*

*Proof.* The case when the modification happened while  $r$  was still the root of its component is trivial since this modification directly changed the version of  $r$ .

The only case left is when the root was another vertex. However, since we know that the second `find_root` invocation returned  $r$ , there should be another modification that made  $r$  the root again — it has to increase the version of  $r$  at the beginning.  $\square$

**Theorem 3.2.** *The connectivity check algorithm in Listing 1 is linearizable.*

*Proof.* Proof by contradiction. There can be only two reasons why the connectivity check is not linearizable: either there was no point in time when  $u$  and  $v$  were in the same connected component, and the check returned `false`, or there was no point in time when  $u$  and  $v$  were in different components, and the check returned `true`.

*Case 1: `connected(u, v)` returned `true`.* Assume that there was no point in time when  $u$  and  $v$  were in the same component. To succeed, all three `find_root` invocations should find the same root  $r$  with the same version.

Consider the connectivity of node  $r$  at the moment immediately after the first root search for  $u$ . According to the assumptions, at this moment  $r$  was either connected to only one of the nodes or not connected to both. In the first case, in order to make another node connected with  $r$  and return it in the `find_root` invocation, there should be at least one edge removal to split the connected node and  $r$  and one edge addition to make a `find_root` invocation for another, disconnected node return  $r$ . The second case is similar: there should be at least two edge additions in the same component of connectivity, the last one of which extends the component of connectivity with  $r$ . As a result, in both cases there were two modifications, and the second one started after the first `find_root` invocation. This way, there

is a modification between the completion of the first root search for  $u$  and the second one, both returned the same root  $r$ . Following the Lemma 3.1, the obtained versions should be different. Contradiction.

*Case 2:  $connected(u, v)$  returned false.* Assume that during this  $connected(u, v)$  invocation  $u$  and  $v$  were always in the same component of connectivity. In this case, the following history should be obtained:

$\mathcal{C}_u$ :	$r_u, V_u$
$\mathcal{C}_v$ :	$r_v, V_v$
$\mathcal{C}_u$ :	$r_u, V_u$
$\mathcal{C}_v$ :	$r_v, V_v$
$\mathcal{C}_u$ :	$r_u, V_u$

Remember that according to the assumption  $u$  and  $v$  have always been in the same component. Therefore, since four root changes of the same component have been detected, there were four corresponding modifications. The third of these modifications started after the first  $connected(u, v)$  invocation and influenced the result of the fourth, thus we can use it to apply the Lemma 3.1 to the first and the last  $connected(u, v)$  invocations and get a contradiction. □

**Other Balanced Trees.** We used the Cartesian trees as an underlying data structure in *ETT*-s because parent links in the Cartesian trees naturally form an acyclic graph at any moment of time and roots do not change due to rebalancing. Otherwise, concurrent readers could have looped endlessly or could have wrongly detected a component change during a modification. In fact, our single-writer multiple-reader implementation of the Cartesian trees is barely changed from the sequential implementation. The main issue with the Cartesian trees is that their time complexity is *expected* due to randomization. So, we need to show how to employ any *deterministic* balanced tree instead, for example, the B-tree [22].

In the B-tree each node contains a number of keys rather than a single key. In our case, a node can be viewed as a container of some vertices and edges from graph  $G$ . As a consequence, to allow root searches we also should store links from vertices to nodes they are contained in. The ideas for atomic merges and splits, and for linearizable queries are the same – in a logical part of modifications one tree is linked/unlinked to/from another tree, each root has a version that changes before a modification. The only problem is that the B-tree can allocate a new node and make it the root during a merge or replace the root during a split, and we want roots not to change unless in the logical part. The solution is straightforward – if the B-tree replaces the root with another node, we can instead swap the content of the root and the content of this node. Note that the structure of the tree remains the same, but concurrent readers now can not see odd root changes.

The number of keys in each B-tree node can be chosen arbitrarily; a B-tree with  $\Theta(k)$  keys in each node supports  $\mathcal{C}_u$  and  $\mathcal{C}_v$  operations in  $\mathcal{O}(k \log_k N)$  time, and  $\mathcal{C}_u$  in  $\mathcal{O}(\log_k N)$ . As a result, if we choose constant  $k$ , we will get the same, but *deterministic* asymptotic time complexities as for the Cartesian tree. What is more,  $\Theta(\log n)$ -ary trees can improve read time complexity for the general dynamic connectivity problem to  $\mathcal{O}(\frac{\log N}{\log \log N})$  as shown by Henzinger and King [1]. Besides, nodes in other balanced binary trees can be viewed as containers of size 1, and thus, we can adapt them in the same way as the B-tree by introducing this additional level of indirection.

## 4 Concurrent Dynamic Connectivity

We design a concurrent generalization for the classical sequential dynamic connectivity algorithm proposed by Holm et al. [2] for several reasons. First of all, this algorithm can be improved to the best known deterministic algorithm [3] by supplementing it with the *lazy local trees* and *shortcut structure*. Given that this algorithm is already complex, it makes sense to start with making a generalization for it. Secondly, it uses the same ideas as the randomized algorithm with amortized time complexity. Moreover, we prioritize average time over worst-case time as it is a better metric for throughput.

## 4.1 The Sequential Algorithm via *ETT*-s

The classic algorithm that solves the dynamic connectivity problem uses Euler Tour Trees to represent different components of connectivity. The idea is to maintain a spanning forest  $F$  of graph  $G$ , so most of the dynamic connectivity operations can be straightforwardly reduced to the *ETT*-s:

- $\text{add}(u, v)$  either merges two *ETT*-s if it connects different components, or adds  $(u, v)$  to the set of non-spanning edges. If the edge  $(u, v)$  connects different components of connectivity, then it should be added to the spanning forest, i.e. Euler Tour Trees.
- $\text{remove}(u, v)$  either removes  $(u, v)$  from the set of non-spanning edges or search for a *replacement* edge – an edge that reconnects the spanning trees after the spanning  $(u, v)$  edge is removed, reconstructing the *ETT* if a replacement is found or splitting it otherwise.
- $\text{find}(u, v)$  finds the roots in the Cartesian trees of vertices  $u$  and  $v$  and compares them.

**The Level Structure.** Note that a full scan of all non-spanning edges in  $\text{add}(u, v)$  leads to a linear time complexity. To solve this problem, Holm et al. introduced a special *level structure* for storing non-spanning edges to search for replacement efficiently.

Each edge  $e$  is assigned with an integer non-negative level  $l(e) \leq l_{max} = \lfloor \log_2 n \rfloor$ . Edges of level  $\geq i$  form the subgraph  $G_i$ . In particular,  $G = G_0 \supseteq G_1 \supseteq \dots \supseteq G_{l_{max}}$ . In each subgraph  $G_i$  a spanning forest  $F_i$  is maintained with Euler Tour Trees.

Furthermore, the level structure supports the following invariants:

- $F = F_0 \supseteq F_1 \supseteq \dots \supseteq F_{l_{max}}$ . Intuitively, this means that to get the spanning forest of  $G_i$  we can take  $F$  and omit edges of the levels lower than  $i$ .
- $F$  is a maximal spanning tree for edge weights defined as their levels.
- Every component of connectivity in  $G_i$  has size at most  $n/2^i$ . As a result, edges cannot have levels greater than  $\lfloor \log_2 n \rfloor$ .

The main idea of how to obtain better amortized operation time is to increase the level of every edge processed during a replacement edge search if it can not be a replacement. As the maximum level is bounded with  $\lfloor \log_2 n \rfloor$ , the total number of level increases is at most  $m \log_2 N$ , where  $m$  is the number of inserted edges.

The  $\text{add}(u, v)$  and  $\text{find}(u, v)$  operations in this scheme work with the lowest level  $G_0 = G$  and are not changed relatively to the general description above. The information about non-spanning edges (the edges that do not belong to the spanning forest) on the level  $l$  is stored in adjacent nodes of the level  $l$ . Specifically, each node of the Cartesian tree has a set of non-spanning edges. The  $\text{remove}(u, v)$  operation when removing a non-spanning tree of level  $l$  just deletes the information about it in the corresponding nodes of  $F_l$ .

Upon a spanning edge of level  $l$  removal, a replacement edge is first searched among the edges of level  $l$ . For this purpose, spanning tree  $T$  is split into trees  $A$  and  $B$  that will appear if there is no replacement edge. Without loss of generality, we can assume that the size of  $A$  is not greater than the size of  $B$ . Then, all spanning edges on the level  $l$  in  $A$  increase their levels by 1. It does not violate the invariant on the sizes of connected components because  $A$  is at least twice less than the component of connectivity  $T$  and the size of  $T$  was at most  $n/2^l$ . After that, the algorithm iterates over all non-spanning edges of level  $l$ . Those edges that can not be a replacement increase their levels as well. They still will be non-spanning in  $G_{l+1}$  due to the previously mentioned spanning edge rises. If an edge can be a replacement, then the search is finished and the edge is added to the spanning forest  $F_i$  for every  $i \leq l$ . If all edges are covered and no replacement is found, the spanning tree  $T$  stays split into  $A$  and  $B$ , and the search continues in the lower levels, until  $l \geq 0$ .

To iterate over all spanning or non-spanning edges of the current level effectively, each node of the Cartesian tree has flags whether there are such edges in its subtree. This way, a new edge on the current level can be found in  $\mathcal{O}(\log N)$  time by descending from the root to children with such edges.

## 4.2 Lock-Free Connectivity Checks

We start constructing our scalable concurrent algorithm for dynamic connectivity iteratively. At first, assume that all edge additions and removals are ordered by a coarse-grained lock or using the flat combining technique.

However, with the lock-free single-writer concurrent *ETT* algorithm presented in Section 3, we immediately get a non-blocking connectivity check for the dynamic connectivity problem by replacing a sequential *ETT* for the lowest



level  $G_0$  with the concurrent one. This replacement is correct since connectivity queries work only with the lowest level of the level structure, while a single writer, in fact, performs all modifications. Besides, if a  $\Theta(\log n)$ -ary balanced tree is employed for the lowest level, query time complexity can be improved to  $\mathcal{O}(\frac{\log n}{\log \log n})$  without affecting average modification time [1].

### 4.3 Fine-Grained Locking for Updates

Intuitively, modifications of distinctive components can be naturally parallelized. That leads to an idea to replace the global lock with fine-grained per-component locking. Nonetheless, a problem has to be solved: the components of connectivity change dynamically, so we need a representative of each component that can hold the lock.

Similarly to the non-blocking reads, connectivity components can be represented by the Cartesian tree roots of the  $ETT$ -s of  $G_0$ . This way, an update operation can obtain the roots via `get_roots` invocations and take the corresponding locks. Once the locks are acquired, we have to check that the found nodes are still roots, i.e the operation locked *some* components. After that, we repeat the root searches for both  $u$  and  $v$ . Since after locking the component of the root can not be modified, when `get_roots` returns the same root, it is guaranteed that the correct representative is acquired, and no operation can modify the corresponding component while we hold the lock.

Listing 2 presents a pseudo-code of the algorithm that locks the required components. The rest of the design stays the same as for the coarse-grained locking.

```

1 func get_roots(u, v) = while (true) {
2   // Get the roots of 'u' and 'v' with the versions
3   (u, l) := get_roots(u, v)
4   (v, r) := get_roots(u, v)
5   // Always have the same lock ordering
6   if u < v swap u, v
7   get_lock(u), get_lock(v)
8   // Re-check that the correct locks were acquired
9   if u == null or v == null or
10    u == v or u == v
11    u != get_roots(u, v) or
12    v != get_roots(u, v)
13    get_unlock(u), get_unlock(v) // unlock
14    continue // and restart
15 return // the locks were taken correctly
16 }
```

**Listing 2:** The per-component fine-grained synchronization for update operations that works with vertices  $u$  and  $v$ . The Cartesian tree roots, returned by `get_roots`, represent the components.

### 4.4 Lock-Free Non-Spanning Edge Updates

Given that at most  $N - 1$  edges can belong to a spanning forest and there can be up to  $\mathcal{O}(N^2)$  different edges in a graph, most modifications in dense graphs do not change the spanning forest. For example, in the real-world hyperlink and physical connections between routers graphs, the ratio of edges over vertices often exceeds 10 [23]. Therefore, about 90% edges of the graph are non-spanning. The optimization proposed in this subsection performs non-spanning edge additions and removals without taking locks.

**High-Level Idea.** The main idea is to assign statuses to the edges and process the ones that are marked as non-spanning in a non-blocking way. Thus, each update operation checks whether the edge is spanning or should become one and processes it under the corresponding locks in this case. Otherwise, the edge is non-spanning, and the operation can be performed in a non-blocking way, re-starting from the beginning if some intermediate check or atomic update fails. The intuition is that such non-blocking operations linearize at the points of status updates.


The principal challenge of non-blocking edge addition lies in a possible conflict of this operation with a concurrent spanning edge removal. If a spanning edge removal searches for replacement and does not find it, while in parallel a non-spanning edge that can be the replacement is added, the data structure state becomes non-linearizable after the component of connectivity split by the removal operation. The conflict indicates the need for an additional channel of communication between spanning edge removals and non-spanning edge additions. To create such a channel,

```

1 class Edge v k a S
2 class Spanned Edge {
3     E = null // 'DONE' when done
4 }
5
6 val = C()
7 var S = null
8
9 func E = while(true) {
10     when {
11         NON-SPANNING -> {
12             // Update the edge status, re-start on failure
13             if !CAS(& NON-SPANNING, REMOVED): continue
14             // Remove the edge from the set of non-spanning
15             E
16             return
17         }
18         SPANNING -> {
19             // Take the lock and remove the edge
20             E
21             return
22         }
23         REMOVED -> return // already removed
24         INITIAL -> return // not added yet
25     }
26 }
27
28 func E = while(true) {
29     if E:
30         // Take the lock and add the edge
31         E
32         return
33     // Try to add a non-spanning edge; status is INITIAL
34     E // add to the multi-set
35     // Is there a concurrent remove for which the edge can be a replacement?
36     E:
37     if E != null and
38         E:
39         if CAS(& E, null, E):
40             // The edge became the replacement
41             E
42             CAS(& INITIAL, SPANNING)
43             return
44         else if E == DONE:
45             E
46             // The edge is about to become spanning, perform the addition under the lock
47             E
48             return
49     // Try to make the edge non-spanning and
50     // check that its endpoints are in the same component
51     if !E:
52         E, continue
53     if CAS(& INITIAL, NON-SPANNING): return // success
54     else:
55         // A concurrent writer uses the edge as a
56         // replacement and moved the status to SPANNING
57         E
58         return
59 }

```

**Listing 3:** Lock-free algorithm for non-spanning edge additions and removals in the simplified case, when all spanning edge updates use coarse-grained locking, non-spanning edges do not organize the level structure and are stored in a single multi-set. In addition, in this code we assume that threads cannot add the same edge concurrently.



**Figure 4:** Possible transitions between edge statuses. Red dashed transitions are performed under the lock.

the removal operation at the beginning publishes information about itself together with an initially empty slot for a replacement edge. This allows a concurrent non-spanning addition to propose its edge as a replacement.

**Assumptions.** The whole algorithm for non-blocking modifications does not fit into space requirements and contains a ton of technical details, so we present a simplified solution that relies on a few assumptions and shows the main idea of the algorithm. At first, we assume that threads cannot add the same edge concurrently. Besides, we offer a solution that uses coarse-grained locking for updates and does not leverage the level structure from the Holm et al. algorithm — we maintain a global set of non-spanning edges instead. We consider applying the presenting approach to the fine-grained locking and the Holm et al. algorithm as a technical detail. However, we fix these assumptions in our implementation, which pseudo-code is presented in Appendix C. Note that we also do not have these assumptions in our experiments.

**Edge Statuses.** Under the discussed assumptions, we need four statuses, the state machine of which is presented in Figure 4:

- `INITIAL` – the initial status assigned to every new edge;
- `SPANNING` – the edge is in the spanning forest, so its removal should be executed under the lock;
- `NON-SPANNING` – the edge is not in the spanning forest, so its removal can be non-blocking;
- `REMOVED` – the edge is removed.

In the complete solution, an additional `ADDING-SPANNING` status is required to synchronize concurrent additions of the same edge. Otherwise, an addition operation can not distinguish cases when an edge was already added to the spanning forest and when it is being added by a concurrent addition, and thus, can not decide whether it should synchronize with a parallel thread to be sure that the edge is added.

**The Data Structure.** Listing 3 presents an algorithm for non-blocking non-spanning edge processing. For simplicity, each edge is wrapped with a special `Edge` class that also stores the status from Figure 4 (line 1). In a practical implementation, one may use a concurrent hash table to store the statuses. What is more, we have `REMOVED` status logically, but in practice, it can correspond to an absence of a value to avoid memory leaks.

Besides, each spanning edge removal operation publishes it to a special global `spanning_edges` field (line 7), so a concurrent edge addition can propose its edge as an replacement to restructure the connectivity component — we use `SpanningEdge` descriptor for this purpose. Finally, all non-spanning edges are stored in a multi-set `non_spanning_edges` (line 6). For fine-grained locking over components of connectivity, `spanning_edges` field and the multi-set should be stored in nodes of the Cartesian trees instead.

**The `remove_edge(u, v)` Operation.** The edge removal procedure (lines 9–26) is straightforward. When the status is already `REMOVED`, then the operation can complete immediately — the edge is already removed (line 23). Similarly, when the status is `INITIAL`, the edge is not added yet, so we can linearize this removal before its addition and finish (line 24). In the case of the `SPANNING` status, the corresponding `ETT` should be modified, so the operation takes the lock and performs this removal as “writer” (line 20). In the last case, the status is `NON-SPANNING`. Thus, changing it atomically to `REMOVED` logically removes the edge (line 13). After that, we delete it physically from `non_spanning_edges` (line 15).

The replacement search, while iterating over the multi-set of non-spanning edges, can see some of them in the `INITIAL` status. For correctness purposes, it should help to add such edges using an algorithm almost identical to the non-spanning edge addition described later.

**The `add_edge(u, v)` Operation.** The edge addition procedure starts with a connectivity check (line 29) and falls to the blocking mode when the edge connects different components (line 31). Then, the algorithm optimistically adds the edge to `edges` (line 34) — it will later remove it from the multi-set if the edge cannot be non-spanning.

Once the edge is added to `edges`, the operation may notice that there is a concurrent spanning edge removal for which this adding edge can be a replacement to reconstruct the Euler Tour Tree (lines 36–38). In this case, we try to set our edge to the `replacement` field of the removal operation (line 39) — remember, the edge status is still `INITIAL`. If the edge is installed to `replacement`, the operation removes the edge from the multi-set (line 41), changes the status to `SPANNING` (line 42) and completes. Note that the status update failure (line 42) can be caused only if the concurrent edge removal already processed the edge as replacement and made it spanning (or even removed after that).

When the edge could be a replacement but the remove operation is logically completed (line 44), then the edge will connect different components. Therefore, we delete it from the multi-set of non-spanning edges and fall back to the blocking mode (lines 45–47).

We successfully examined the case when a concurrent removal operation has been detected. Since the edge is already added to `edges`, if a concurrent removal starts after we read `replacement` from `edges`, it should find our edge. However, we still need to re-check that the edge does not connect different components, removing the edge from `edges` and re-starting the operation in this case (line 51–52).

Finally, we can add the edge as a non-spanning by an atomic status update from `INITIAL` to `NONSPANNING` (line 53). If the corresponding CAS succeeds, the edge is successfully added as a non-spanning, so the operation completes. Otherwise, if the CAS fails, a concurrent “writer” either has used the edge as a replacement to reconstruct the *ETT* or has helped to add the edge, so one copy of it should be removed from the multi-set of non-spanning edges, and the operation completes (lines 55–57).

**Correctness of Non-Spanning Edge Addition.** To show the correctness of the described non-spanning edge modification, we prove the lemma below, which shows the correctness of non-spanning edge addition with a concurrent edge removal. We consider the operations under locks and isolated non-spanning edge modifications trivial, so the corresponded proofs are omitted.

**Theorem 4.1.** *In the proposed non-blocking non-spanning edge addition algorithm (Listing 3), the conflict with a concurrent spanning edge removal that leads to a non-spanning edge between different components of connectivity is impossible.*

*Proof.* Proof by contradiction. First, if there is a replacement for the removed edge, then the concurrent addition does not play any role. Similarly, if there is no replacement and the edge of the concurrent addition cannot be a replacement, the conflict is impossible. Therefore, we consider the situation when the non-spanning edge addition provides the only potential replacement for the concurrent removal.

If the non-spanning addition inserts the edge to `edges` (line 34) before the removal operation starts its blocking part, then the removal operation will see the edge and help complete its addition, so there will be no conflict. On the other side, if the removal operation finishes before the connectivity check in the addition algorithm (line 53), this check will fail and cause the operation to fall back to the blocking version, so the conflict will not arise. Therefore, the only interesting case is when the removal operation still holds the lock, while the non-blocking addition is between the lines 37-47. As a consequence, the `edges` invocation on the line 38 will return correct results if a conflict with a removal is possible. Consider what happens first: the removal operation publishes itself in the `edges` or the non-spanning addition operation reads `edges` (line 36).

*Case 1:* the publication is the first, so the addition will see the removal operation and will suggest the edge as a replacement instead of completion of the addition. If the replacement proposal (CAS) fails, it means that there is another edge in the slot or the removal closed the slot; both situations are correct.

*Case 2:* the read is the first, consequently, the information of the edge is inserted before the search, because the publication precedes the search and the information insertion precedes the read (line 34). Thus, the search will see the edge and will help to add it if necessary.  $\square$

**Implementation Details.** The discussed simplified version can be adapted to the Holm et al. algorithm that we use as a sequential version under the hood. Due to space limit and dramatic complexity increase, we move the pseudo-code of our full algorithm with all the implementation details to Appendix C.

## 4.5 Time Complexity Analysis

Consider an abstract fully lock-free generalization of the Holm et al. algorithm for the dynamic connectivity problem, e.g., the one implemented via lock-free software transactions. Since the replacement search takes  $\Omega(N)$  in the worst case, and other operations can make it fail and re-run, the tight upper bound on the total work is  $\Omega(MN)$ , which is substantially worse than  $\mathcal{O}(M \log^2 N)$  of the sequential algorithm. This observation motivated us to design spanning edge updates blocking rather than lock-free.

Theorem 4.2 below shows the total work performed by all threads *excluding the work spent on waiting for locks* for our full algorithm. The same bound can be used when some of the optimizations (such as non-blocking modifications) are switched off since they mainly reduce the need for locks, which is not essential in the analysis.

**Theorem 4.2.** *The total amount of work excluding the work spent on locking is  $\mathcal{O}(M \log^2 N + P \cdot M \log N)$ , where  $P$  is the number of processes,  $M$  is the total number of operations, and  $N$  is the number of graph vertices.*

*Proof.* Obviously, for  $P = 1$  our algorithm performs the same amount of work as the sequential algorithm by Holm et al. To prove the theorem in the concurrent scenario, we need to show that each successful modification can cause at most  $\mathcal{O}(\log N)$  additional work for every other thread.

- A successful spanning forest modification induces  $\mathcal{O}(\log N)$  extra work for concurrent connectivity queries, non-spanning additions, or non-spanning removals due to re-tries.
- Non-blocking additions and removals cause at most  $\mathcal{O}(1)$  work for other modifications due to edge state changes and at most  $\mathcal{O}(\log N)$  work for a replacement edge search.
- Another  $\mathcal{O}(\log N)$  extra work is performed by a blocking removal to find a replacement in case a concurrent addition has already proposed a replacement.

□

## 5 Evaluation

To evaluate our concurrent optimizations and their combinations we perform a range of experiments on workloads composed of both real and synthetic graphs. We consider scenarios with different ratios of reads to updates, which help us evaluate the non-blocking reads, as well as the applicability of our algorithms to different workloads. In addition, we consider scenarios with only additions or only deletions to compare the performance of these two operations. All scenarios are executed on graphs with different properties, e.g. with different graph densities. Our experimental setup is publicly accessible [24].

### 5.1 Evaluation Setup

**Benchmarks.** We use three different workloads in the experiments:

- **The Random Subset Scenario.** Here, the data structure is initially filled with half of the graph chosen randomly. After that, multiple parallel threads execute random operations on randomly chosen graph edges. To maintain the number of edges in the structure approximately constant, the percentage of edge additions is equal to the percentage of edge removals, which means that only the ratio of reads to modifications can change. Aksenov et al. [8] essentially use the same experimental setup.
- **Incremental Scenario.** In this case, threads add concurrently the whole graph to the initially empty data structure. This benchmark is inspired by the *incremental dynamic connectivity problem*.
- **The Decremental Scenario.** This benchmark is opposite to the previous one: parallel threads remove graph edges from a data structure initialized with the whole graph.

**Graphs.** We used both real graphs and random graphs listed in Tables 1 and 2. We divide these graphs into relatively small and large. Smaller graphs help to analyze the dependency between the number of threads and the algorithm performance. Large graphs are always processed by the maximum number of threads. The graph of USA roads was taken from a competition on finding shortest paths [25]. Other real-world graphs are from the SNAP graph repository [26]. The random graphs were generated by the Erdős-Rényi model, where all edges are present with an

equal probability. Large synthetic graphs are from another edition of the DIMACS graph processing competition [27]. We remove loops and multi-edges from the graphs, as they do not affect connectivity.

**Table 1:** Small graphs used in our experiments for evaluation all the algorithms on various numbers of threads.

Graph	V	E	Description
USA roads	435666	521200	Colorado State roads
Twitter	81306	1342296	Social circles in Twitter
Stanford web	281903	1992636	Web graph of Stanford.edu
Random, $ E  =  V $	400000	400000	Random sparse graph
Random, $ E  = 2 V $	300000	600000	Random sparse graph
Random, $ E  =  V  \log  V $	100000	1600000	Random dense graph
Random, $ E  =  V  \sqrt{ V }$	20000	1600000	Random high-density graph
Random, 10 components	100000	1600000	Random graph with 10 components

**Table 2:** Large graphs used in our experiments for evaluating all the algorithms on the maximum (144) threads. As the experiments on large graphs are time-consuming, and the trends seem similar to the ones on smaller graphs (listed in Table 1), we find it sufficient to have experiments only on the maximum number of threads.

Graph	V	E	Description
Full USA roads	23.9M	28.9M	The graph of all roads in the USA
LiveJournal	4.8M	42.9M	Friendship graph in LiveJournal
Kron	2.1M	91M	Synthetic Kronecker graph
Random	4.2M	48M	Erdős-Rényi random graph

**Hardware and Software.** All the algorithms presented in this paper are implemented in Kotlin for JVM. The experiments were run on a server with 4 Intel Xeon Gold 6150 (Skylake) sockets; each socket has 18 2.70 GHz cores, each of which multiplexes 2 hardware threads, for a total of 144 hardware threads. We used OpenJDK 11<sup>1</sup> in all the experiments and the Java Microbenchmark Harness (JMH) for running benchmarks [28].


## 5.2 Algorithm Combinations

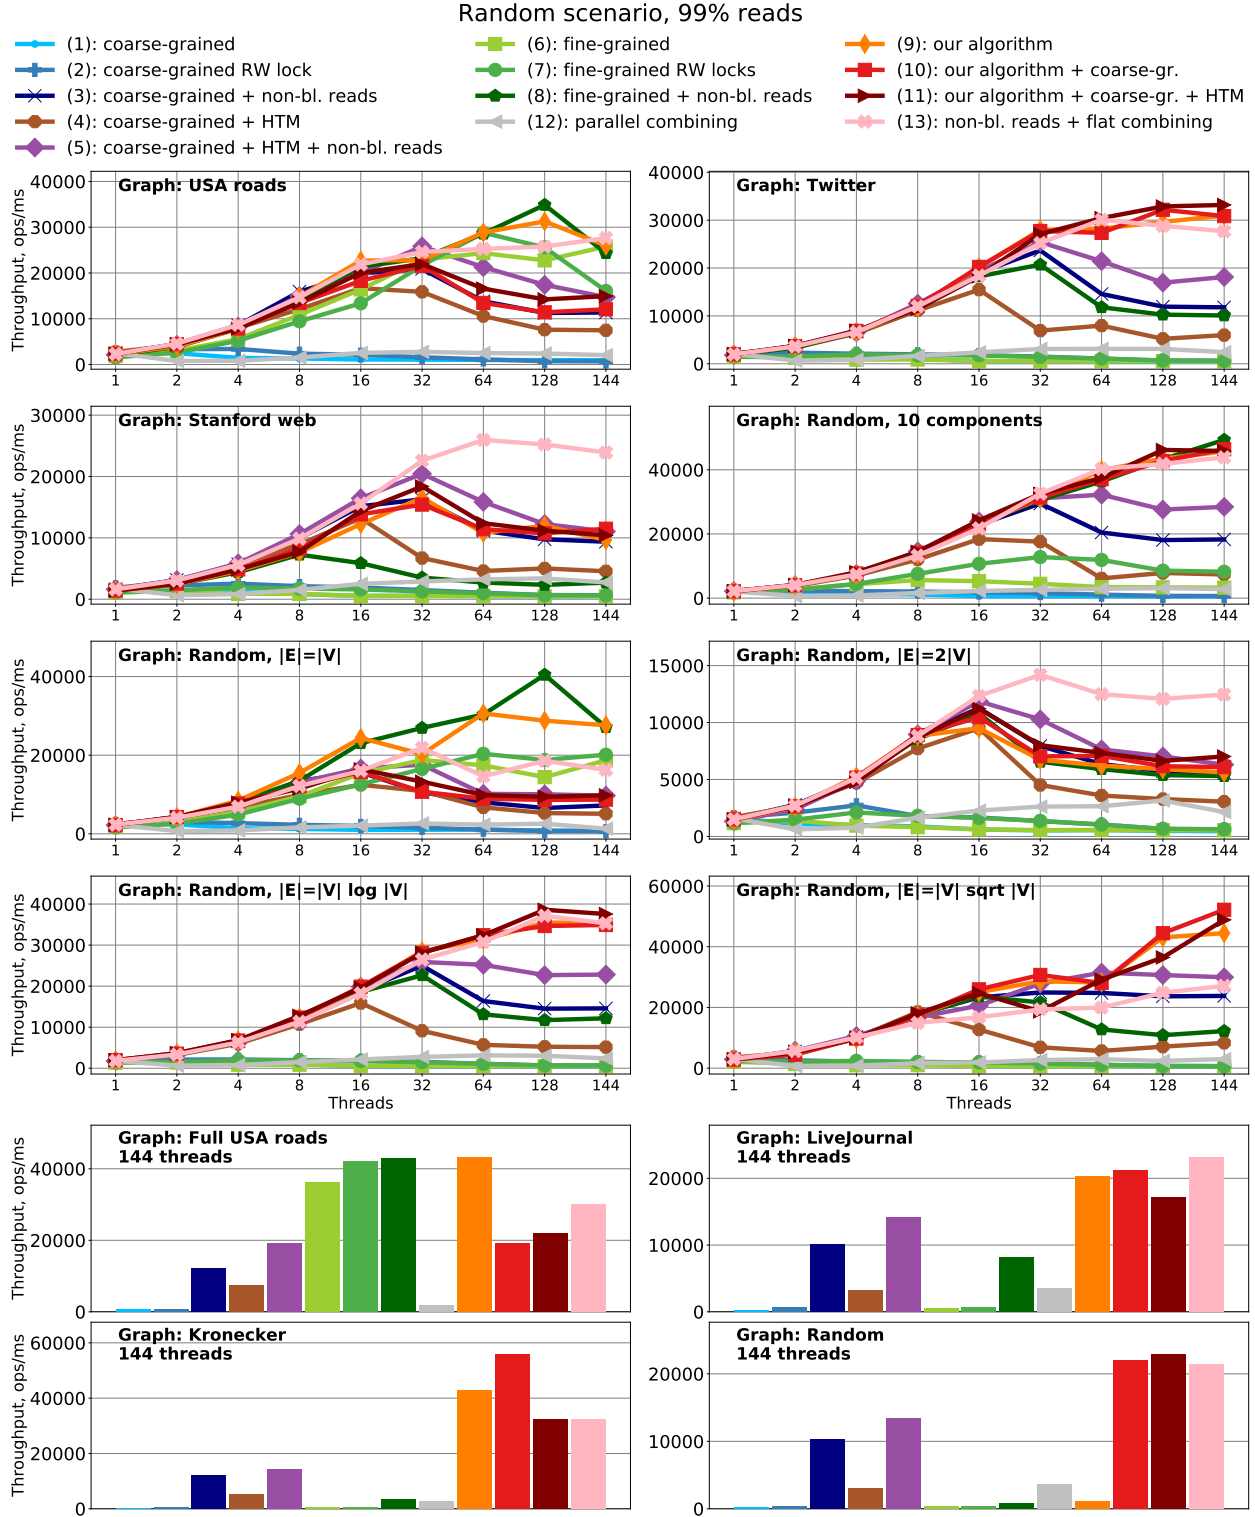
Our experiments examine a wide variety of techniques and scenarios. We enumerate them to make our plots easier to read. First, (1) coarse-grained locking can be used to synchronize *all* operations. We can improve it by (2) using a readers-writer lock instead, (3) using the non-blocking `AtomicMarkableReference` implementation, leveraging hardware transactions (HTM) via lock elision [29, 30]. This can be done for (4) all operations, or (5) only for updates. Similarly, (6) the solution with fine-grained locks can be improved with (7) readers-writer locks and with (8) non-blocking `AtomicMarkableReference`. In addition, the algorithm for non-blocking non-spanning edge updates can be applied to the versions with both (9) fine-grained and (10) coarse-grained locking, each uses non-blocking `AtomicMarkableReference` in our experiments. It is also possible to use non-blocking non-spanning edge updates and `AtomicMarkableReference`-s with HTM (11) for spanning edge additions and removals. In addition, our optimizations are compared with (12) the *parallel combining* technique [8], which processes read operations in parallel, and (13) a similar *flat combining* algorithm, but leveraging our non-blocking reads. We examine the performance of all these variants in the following, keeping the numbering consistent.

**Sampling.** Iyer et al. discovered that for better performance, dynamic connectivity algorithms should use *random sampling* even in the deterministic Holm et al. algorithm [7]. This random sampling is employed as a fast path in the replacement search to prevent the expensive promotion of spanning edges to the next level. We see the same tendency in the concurrent scenario, so we use the sampling heuristic for all algorithms in our experiments. Furthermore, it is even more important in the concurrent case as it reduces the average time when spanning edge removals hold the locks.

<sup>1</sup>We use Java 11 because `AtomicMarkableReference` does not work with the recent Java versions, but this should not affect the analysis significantly.

does not work with the recent Java versions, but this should not affect the





**Figure 6:** Experimental results on the *random* scenario with 80% connectivity checks and 1% updates. **Higher is better.** The top plots show the average operation throughput on small graphs, listed in Table 1, while the bar charts below show the results on large graphs, listed in Table 2, on maximum parallelism.



### 5.3 Experimental Results

**Lock-Free Reads.** The first observation we make is that our non-blocking connectivity query is consistently the most efficient way to implement the  $\mathfrak{a}$  operation. The readers-writer lock does not scale even for read-dominated scenarios and, in fact, is worse than normal locks. Flat combining with non-blocking reads performs better than the read-optimized flat parallelization technique. Moreover, our non-blocking reads do not have worse single-threaded performance, since the additional memory locations they read are likely to be in the cache. What is more, in fact, more than 99.99% reads succeed on the first try, so there is no thread starving and they are *practically* wait-free.

**Fine-Grained Locking.** The benefits of using fine-grained locking over connected components are less clear-cut. For graphs that have multiple connected components, the algorithm scales substantially better, but for graphs with a single component, it is slower than coarse-grained locking. Surprisingly, there is an exception to this rule. In the incremental and decremental scenarios for the USA roads graph, which has only one connected component, the algorithms with fine-grained locking scale well. We believe that this is the consequence of the fact that the road graph is planar – when a small number of random edges are removed from a planar graph, the graph quickly loses its connectivity.

**Hardware Transactional Memory.** The algorithms with the lock elision technique have better performance than the ones with coarse-grained locking on our random scenario, except for our full algorithm, where the performances match. However, they are worse in modification-only scenarios, because due to contention hardware transactions are likely to be aborted. For our algorithm there is no difference between coarse-grained locking and the lock elision technique, since the goal of our algorithm is to minimize the number of lock acquisitions.

**Non-Blocking Modifications.** The algorithm variant with non-blocking modifications works well for dense graphs, such as the Twitter graph. It can scale in scenarios with many modifications. However, the trade-off is that this algorithm is slower for single-threaded computations. The only detrimental case is the decremental scenario. The reason for the performance drop lies in the fact that spanning edge removals are far slower than non-spanning ones ( $\mathcal{O}(\log^2 N)$  vs  $\mathcal{O}(1)$ ), so not much work can be parallelized.

**Flat combining.** In most scenarios, the fastest algorithm is either our algorithm with coarse-grained or fine-grained locking, or a flat combining algorithm with our non-blocking reads. Specifically, our algorithm is scalable in many practical scenarios, where queries dominate and updates are relatively infrequent. However, we also emphasize that in “bad” scenarios, where there are many spanning edge modifications and a single connected component, the flat combining algorithm is better due to low synchronization costs and better cache locality.

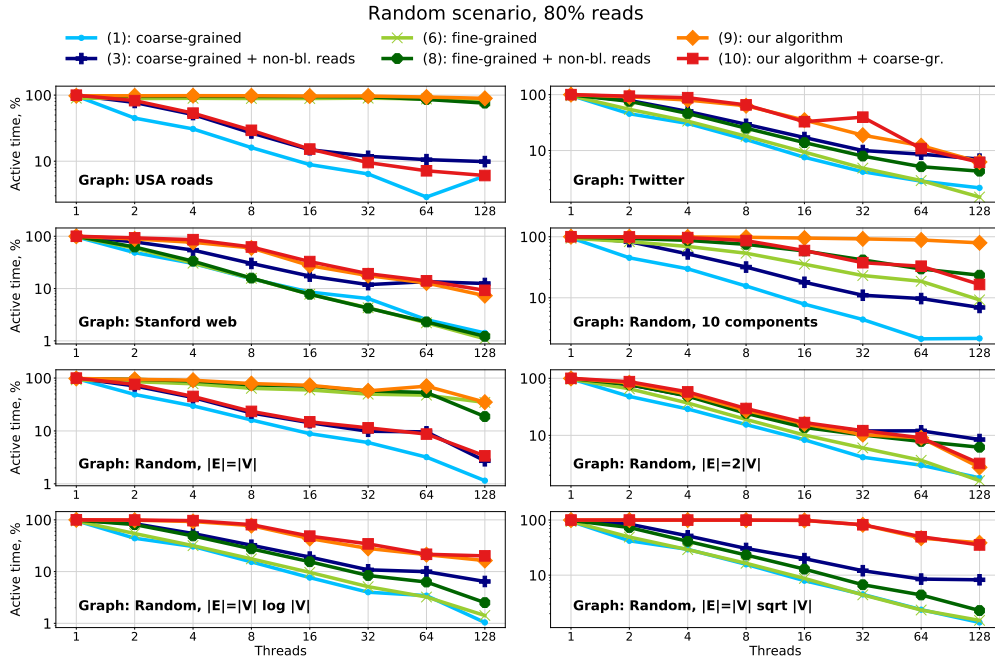
**Non-Spanning Edge Modifications.** Table 3 shows the rates of non-spanning edge additions and removals along with the size of the largest connectivity component obtained during the experiments under the *random scenario*. All the data were obtained on small graphs, listed in Table 1, and on the sequential workload. Our internal experiments show that the collected statistics do not change with increasing the number of threads. Here we see that on dense graphs (such as Twitter, Stanford web, and random graphs with  $|V| \log |V|$  and  $|V| \sqrt{|V|}$  edges) more than 94% of additions and 74% of removals do not require locks. As for the road and sparse random graphs, the optimization for non-spanning edge modifications does not play a significant role — this matches with the experimental results.

**Table 3:** The *random scenario* statistics: shows the rates of non-spanning edge additions and removals along with the size of the largest connected component (divided by the total number of nodes) discovered during the experiment. The statistics have been collected on the sequential workload.


Graph	% of non-span. additions	% of non-span. removals	Largest component, %
USA roads	6.3	1.5	0.4
Twitter	98.7	88.6	94.2
Stanford web	94.4	74.5	74.0
Random, $ E  =  V $	0.1	0.0	0.9
Random, $ E  = 2 V $	63.4	16.0	77.8
Random, $ E  =  V  \log  V $	100.0	87.5	100.0
Random, $ E  =  V  \sqrt{ V }$	100.0	97.5	100.0
Random, 10 components	100.0	87.4	10.0

**Lock Waiting Expenses.** To show how much time various algorithms wait for the required locks, we collect the statistics on the *active* time rate, which excludes the waiting part. In particular, when the rate is 100%, the algorithm

spends all the time in the graph processing and no time on waiting for locks. The corresponding data is presented in Figures 7–8. First, fine-grained locking significantly improves the straightforward coarse-grained implementation on sparse graphs but shows similar performance on the dense ones. At the same time, applying non-blocking reads always reduces the time spent on waiting for locks — this is expected since we use read-dominated scenarios in our experiments. As for the non-spanning edge updates optimization, the impact correlates with the already discussed statistics presented in Table 3 — the lock waiting time is significantly reduced on dense graphs.



**Figure 7:** The *active time* rate for the *random* scenario with 80% connectivity checks and 20% updates that shows the graph processing time and excludes the time spent on waiting for locks. **Higher is better, 100% is the best possible.**



**Figure 8:** The *active time* rate for the *random* scenario with 99% connectivity checks and 1% updates that shows the graph processing time and excludes the time spent on waiting for locks. **Higher is better, 100% is the best possible.**

**Incremental and Decremental Scenarios.** We omit the experimental data for the *incremental* and *decremental* scenarios in the main body and present them in Appendix B.

## 6 Conclusion

We presented a first concurrent variant of the dynamic connectivity, based on the classic sequential algorithm of Holm et al. [2] and a concurrent single-writer Euler Tour Tree data structure. We have investigated a wide range of technical options for implementing our concurrent design, with various trade-offs. While we are not able to identify a single “best-performing” algorithm across all scenarios we investigated, our experimental validation does identify several rules-of-thumb, which should be useful to guide practical implementations. In terms of future work, we plan to investigate NUMA-aware and recoverable versions of these algorithms.

## References

- [1] Monika R. Henzinger and Valerie King. Maintaining minimum spanning trees in dynamic graphs. In Pierpaolo Degano, Roberto Gorrieri, and Alberto Marchetti-Spaccamela, editors, *Automata, Languages and Programming*, pages 594–604, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg.
- [2] Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48(4):723–760, July 2001.
- [3] Christian Wulff-Nilsen. Faster deterministic fully-dynamic graph connectivity. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’13, page 1757–1769, USA, 2013. Society for Industrial and Applied Mathematics.
- [4] Shang-En Huang, Dawei Huang, Tsvi Kopelowitz, and Seth Pettie. Fully dynamic connectivity in  $o(\log n(\log \log n)^2)$  amortized expected time. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’17, page 510–520, USA, 2017. Society for Industrial and Applied Mathematics.
- [5] Mihai Patrascu and Erik Demaine. Logarithmic lower bounds in the cell-probe model. *SIAM Journal on Computing*, 35, 03 2005.
- [6] Mihai Patrascu and Mikkel Thorup. Don’t rush into a union: Take time to find your roots. *Computing Research Repository - CORR*, 02 2011.
- [7] Raj Iyer, David Karger, Hariharan Rahul, and Mikkel Thorup. An experimental study of poly-logarithmic fully-dynamic connectivity algorithms. *Journal of Experimental Algorithmics*, 02 2000.
- [8] Vitaly Aksenov, Petr Kuznetsov, and Anatoly Shalyto. Parallel combining: Benefits of explicit synchronization. In *OPODIS*, 2018.
- [9] Danny Hendler, Itai Incze, Nir Shavit, and Moran Tzafir. Flat combining and the synchronization-parallelism tradeoff. In *Proceedings of the Twenty-Second Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA ’10, page 355–364, New York, NY, USA, 2010. Association for Computing Machinery.
- [10] Umut A. Acar, Daniel Anderson, Guy E. Blelloch, and Laxman Dhulipala. Parallel batch-dynamic graph connectivity. In *The 31st ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA ’19, page 381–392, New York, NY, USA, 2019. Association for Computing Machinery.
- [11] Maurice P. Herlihy and Jeannette M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, July 1990.
- [12] Jean Vuillemin. A unifying look at data structures. *Commun. ACM*, 23(4):229–239, April 1980.
- [13] Alexander Fedorov, Nikita Koval, and Dan Alistarh. A scalable concurrent algorithm for dynamic connectivity. In *The 33th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA ’21. Association for Computing Machinery, 2021.
- [14] Robert Endre Tarjan. Efficiency of a good but not linear set union algorithm. *Journal of the ACM (JACM)*, 22(2):215–225, 1975.
- [15] 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.
- [16] Yossi Shiloach and Shimon Even. An on-line edge-deletion problem. *Journal of the ACM (JACM)*, 28(1):1–4, 1981.

- [17] Mikkel Thorup. Near-optimal fully-dynamic graph connectivity. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, STOC '00, page 343–350, New York, NY, USA, 2000. Association for Computing Machinery.
- [18] Zhengyu Wang. An improved randomized data structure for dynamic graph connectivity. *ArXiv*, abs/1510.04590, 2015.
- [19] Bapi Chatterjee, Sathya Peri, Muktikanta Sa, and Nandini Singhal. A simple and practical concurrent non-blocking unbounded graph with linearizable reachability queries. In *Proceedings of the 20th International Conference on Distributed Computing and Networking*, ICDCN '19, page 168–177, New York, NY, USA, 2019. Association for Computing Machinery.
- [20] Tom Tseng, L. Dhulipala, and G. Blelloch. Batch-parallel euler tour trees. In *ALENEX*, 2019.
- [21] L. Dhulipala, Julian Shun, and G. Blelloch. Low-latency graph streaming using compressed purely-functional trees. *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2019.
- [22] R. Bayer and E. McCreight. Organization and maintenance of large ordered indices. In *Proceedings of the 1970 ACM SIGFIDET (Now SIGMOD) Workshop on Data Description, Access and Control*, SIGFIDET '70, page 107–141, New York, NY, USA, 1970. Association for Computing Machinery.
- [23] Guy Melançon. Just how dense are dense graphs in the real world? a methodological note. pages 1–7, 05 2006.
- [24] Our experiments on the concurrent dynamic connectivity. <https://arxiv.org/abs/2105.08441>, 2021.
- [25] 9th DIMACS Implementation Challenge - Shortest Paths. <https://dimacs.rti.edu/Challenges/9/ShortestPaths/>, 2006.
- [26] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/>, June 2014.
- [27] 10th DIMACS Implementation Challenge - Graph partitioning and graph clustering. <https://dimacs.rti.edu/Challenges/10/GraphPartitioning/>, 2012.
- [28] JMH - Java Microbenchmark Harness. <https://openjdk.org/projects/comp/benchmarks/jmh/>, 2021.
- [29] Maurice Herlihy and J. Eliot B. Moss. Transactional memory: Architectural support for lock-free data structures. *SIGARCH Comput. Archit. News*, 21(2):289–300, May 1993.
- [30] 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.

## A Non-Linearizability Problem when Omitting the Last Root Search in Listing 1

**Theorem A.1.** *Omitting the last root search in the algorithm from Listing 1 (line 12) results in a non-linearizable connectivity check.*

For readability, we repeat the algorithm for single-writer Euler Tour Trees.


```

1 func findRoot: Edges while (true) {
2   // Get the roots of 'u' and 'v' with the versions
3   (u_root := findRoot(u, v))
4   (v_root := findRoot(v, u))
5   // Has the component of 'u' been changed?
6   if (u_root != findRoot(u, v)): continue
7   if (v_root != findRoot(v, u)): continue
8   // 'u' and 'v' are likely to be in different
9   // components, re-check that 'u_root' and
10  // 'v_root' were snapshotted atomically
11  if (u_root != findRoot(u, v)): continue
12  if (v_root != findRoot(v, u)): continue
13  return (u_root == v_root)
14 }
15 // Returns the root with the current version
16 func findRoot(u, v): (Edges, Edges) {
17   cur := u
18   while (true):
19     p := findParent(cur) // read the parent of 'cur'
20     if p == null: break // is 'cur' the root?
21     cur = p // go towards the root
22   // Return the found root with the current version
23   return (cur, cur)
24 }

```

*Proof.* We show a counter-example for the algorithm without the last root search.

Consider four nodes:  $u, v, w, r$ . Assume that they are connected with the following parent links:  $u \rightarrow w, v \rightarrow w, w \rightarrow r$ . Figure 8 illustrates the structure.



**Figure 8:** Illustration for the counter-example to the connectivity check without the last root search.

A reader checks connectivity of  $u$  and  $v$ . Initially, all versions are 0. The following interleaving with one reader (**R**) and one writer (**W**) leads to a non-linearizable result:

1. **R:** The reader's first root search starts and reaches vertex  $r$ .
2. **W:** A concurrent writer removes the edge  $(w, r)$ ; thus, removing the parent link  $w \rightarrow r$ . The versions of  $r$  and  $w$  are increased by 1.
3. **W:** After that, the writer starts an addition of the previously removed edge, but updates only versions of  $r$  and  $w$  at the current moment. Therefore, the versions of  $r$  and  $w$  become 2.
4. **R:** The first root search finishes with the result  $(r, 2)$ .
5. **R:** The second search starts, reaches  $w$ , finds out that it has no parent, and becomes preempted by the scheduler.

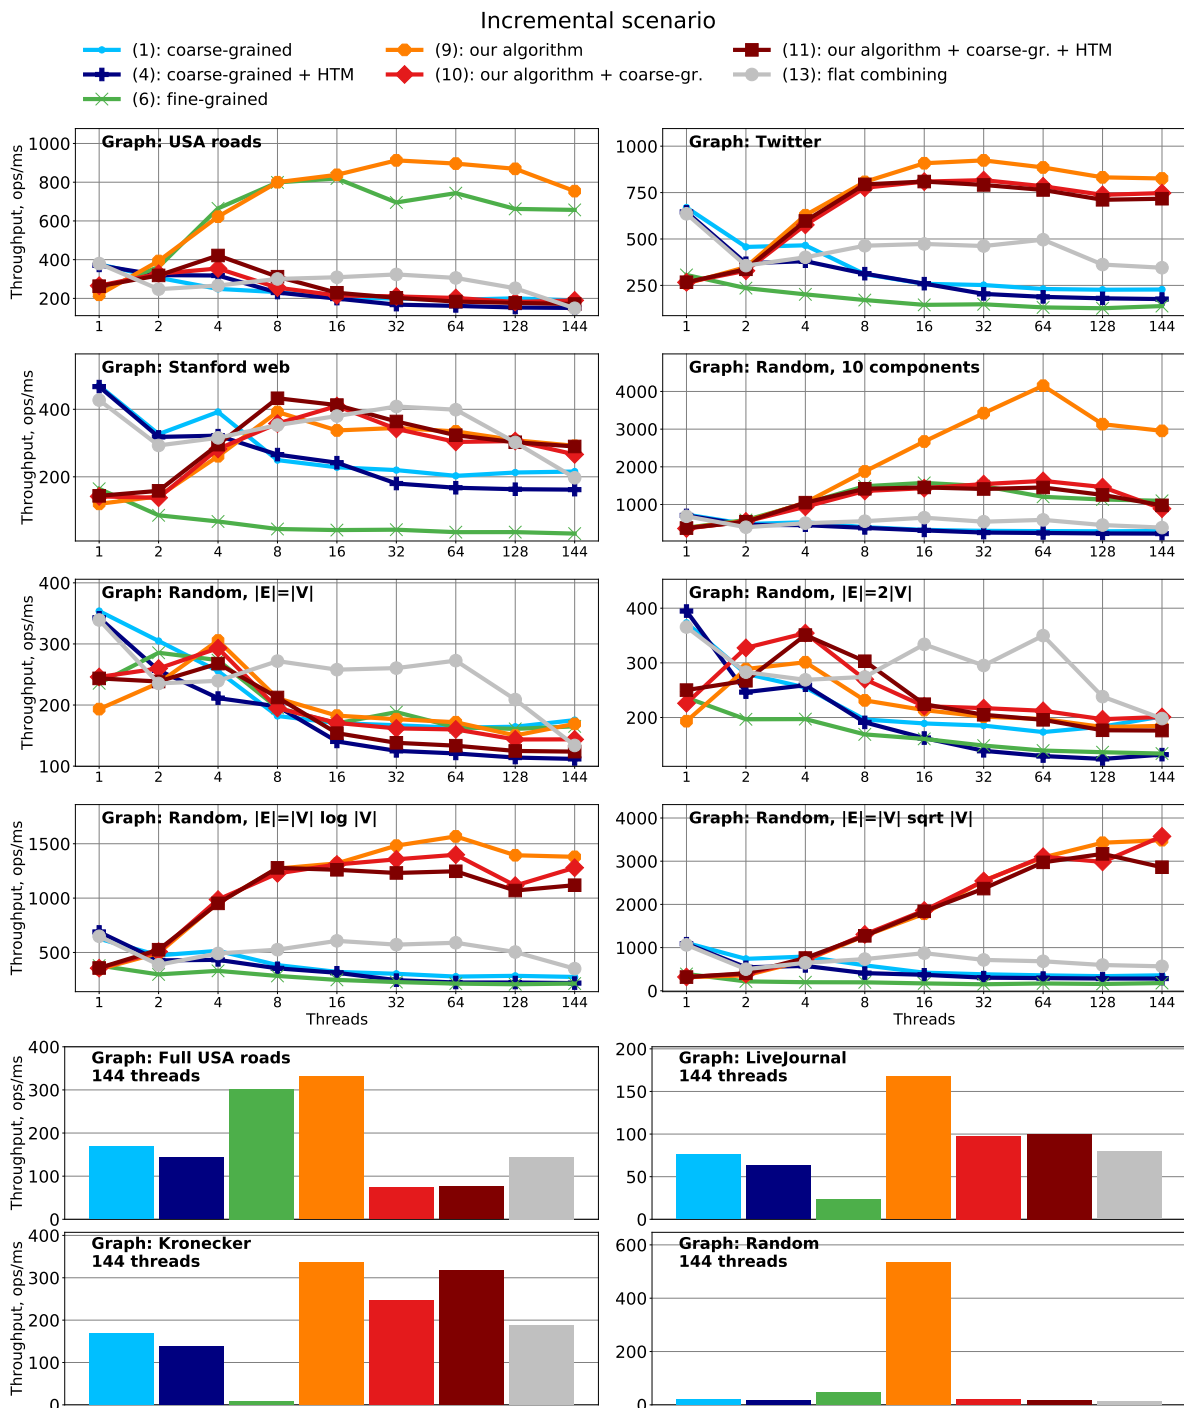
6. **W:** The writer finishes the addition of edge  $(w, r)$ , adding parent link  $r \rightarrow w$ .
7. **W:** The writer starts removing edge  $(w, r)$  again. It increments the version of  $w$  at first.
8. **R:** The reader reads the version of  $w$  equal to 3 and finishes the root search.
9. **R:** The reader starts the third root search and finds the root  $r$  with version 2.
10. **W:** The writer finishes the removal.
11. **R:** The fourth root search finds root  $w$  with version 3.

The vertices  $u$  and  $v$  never were in different components. However, the algorithm without the last root search returns  $\bar{u}$ .

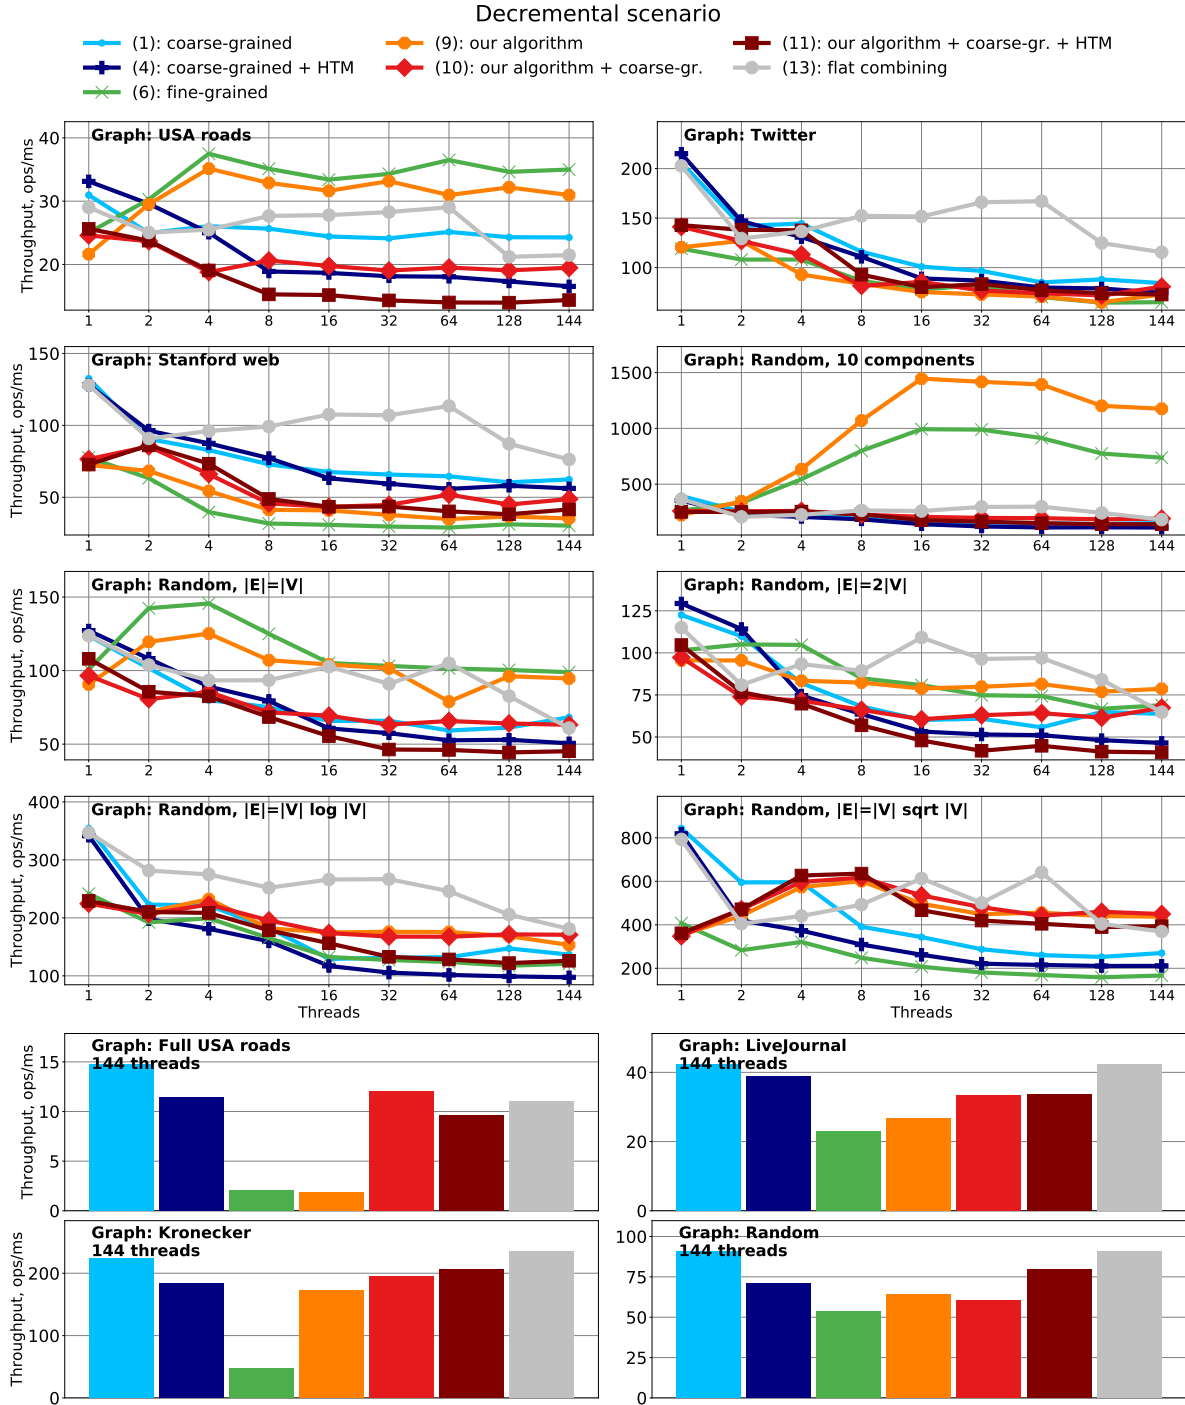
algorithm without the

□

## B Evaluation on the Incremental and Decremental Scenarios




**Figure 9:** Experimental results on the *incremental* scenario. **Higher is better.** The top plots show the average operation throughput on small graphs, listed in Table 1, while the bar charts below show the results on large graphs, listed in Table 2, on the maximum number of threads.




**Figure 10:** Experimental results on the *decremental* scenario. **Higher is better.** The top plots show the average operation throughput on small graphs, listed in Table 1, while the bar charts below show the results on large graphs, listed in Table 2, on the maximum number of threads.





**Figure 11:** The *active time* rate for the *incremental* scenario that shows the graph processing time and excludes the time spent on waiting for locks. **Higher is better, 100% is the best possible.**



**Figure 12:** The *active time* rate for the *decremental* scenario that shows the graph processing time and excludes the time spent on waiting for locks. **Higher is better, 100% is the best possible.**

**Table 4:** The statistics for *incremental* (on the left) and *decremental* (on the right) scenarios, which shows the rates of non-spanning edge additions/removals. The statistics have been collected on the sequential workload.

(a) Incremental scenario

Graph	% of non-spanning additions
USA roads	3.7
Twitter	81.3
Stanford web	66.8
Random, $ E  =  V $	2.9
Random, $ E  = 2 V $	19.4
Random, $ E  =  V  \log  V $	76.9
Random, $ E  =  V  \sqrt{ V }$	93.4
Random, 10 components	76.9


(b) Decremental scenario

Graph	% of non-spanning removals
USA roads	3.7
Twitter	81.3
Stanford web	66.8
Random, $ E  =  V $	2.9
Random, $ E  = 2 V $	19.4
Random, $ E  =  V  \log  V $	76.9
Random, $ E  =  V  \sqrt{ V }$	93.4
Random, 10 components	76.9

## C Lock-Free Non-Spanning Edge Updates: Full Algorithm and Implementation Details

In Subsection 4.4 of the main body, we presented lock-free non-spanning edge updates in the simplified case when concurrent additions of the same edge are prohibited, there is a *global* set of non-spanning edges, and the algorithm uses coarse-grained locking. In this appendix section we show how to get rid of these simplifications.

**Edge Statuses.** The main difference in the state machine is that we need an additional status `IN_PROGRESS` for spanning edge additions. Previously, we changed the edge status at the beginning of spanning edge addition to indicate concurrent removals that they should take locks to proceed. However, concurrent additions of the same edge can not distinguish the following two cases: the case when the edge is already presented and nothing should be done, and the case when the edge is being added concurrently and the operation should wait until its addition is finalized. To fix the problem, at the beginning of spanning edge addition the edge status becomes `IN_PROGRESS` and is changed to `SPANNING` only after all necessary ETT modifications. The updated state machine is presented in Figure 13.



**Figure 13:** Possible transitions of edge statuses in the full version of the algorithm, without assumptions. Red dashed transitions are performed under locks.

**The Data Structure.** Listing 5 shows data structures used in our algorithm and their interfaces. The main class includes ETT levels (line 3) and a concurrent hash map for edge statuses and levels. The `REMOVED` status corresponds to an absence in this hash table. An edge status and a level are stored together because it allows to change them atomically, and as a result, simplifies memory management. For performance purposes, an edge level and a status can be merged to fit in a machine word, since they require only  $\mathcal{O}(\log \log N)$  and  $\mathcal{O}(1)$  bits respectively. In the case of this optimization, to avoid the ABA problem we pair `INITIAL` status with random bits, this way an edge can not be added twice if other threads can help.

As for ETT Nodes, each node contains, in addition to pointers to other nodes, a version used for non-blocking reads, a multi-set with non-spanning edges adjacent to the corresponding vertex, flags needed for efficient navigation in the tree, a field for announcing spanning edge removals, and a lock. All these fields, except for the version, either were not present in the simplified case or were global for the whole data structure.

**Edge Management.** For storing non-spanning edges in Cartesian tree nodes our algorithm uses a concurrent lock-free multiset, which allows iterating over its elements. It is a multiset rather than a set because we permit multiple copies of the same edge to be present simultaneously. The invariant that we maintain is that if a non-spanning edge of level  $r$  is in the graph, then it should have at least one copy in multisets in corresponding adjacent nodes of  $F_r$ . This invariant is easy to provably satisfy: an operation that tries to add a non-spanning edge inserts information about it *before* the linearization point, a removal operation deletes the information only *after* the linearization point. Similarly, when raising an level  $r$  edge, we firstly add the information to level  $r + 1$  and only after the CAS transition from `NON-SPANNING, r` to `NON-SPANNING, r+1` delete the information from the previous level. If some inserted information becomes “useless”, e.g. if the CAS failed, the same operation deletes it to prevent memory leaks.

A more complex problem is that the Cartesian trees should have correct flags “whether there are non-spanning edges in the subtree of a node” in order to be able to iterate over all current level non-spanning edges, spending only  $\mathcal{O}(\log N)$  work per edge. In fact, it will not break the correctness of the algorithm if a small number of nodes has these flags falsely set to true. A replacement search will traverse some odd nodes, but it can potentially fix their flags via backward recursion. So, the idea is to set all flags to `true` in the path between a node and the root when adding information about a non-spanning edge to the node, but in the same time not to change any flags when removing the information. In other words, flags can be set to `true` only by a replacement search *under the locks*. However, the replacement search can mistakenly set a flag to `true`, when there is a concurrent addition that has just added information about a new edge and set the flag to `true`. Our observation is that a simple re-check after setting to `true` (lines 33-37 in Listing 6) can help to avoid incorrect flags (Lemma C.1).

```

1 class CTreeNode In {
2   // Spanning forests for each graph level
3   val F = CTreeNode(h + 1)
4   // Level and status for each edge
5   val e = CTreeNode.S(e)
6   func e()
7   func e()
8   func e(): Bb
9 }
10
11 class CTreeNode {
12   func e()
13   func e()
14   func e(): Bb
15   // Returns the node corresponding to a vertex
16   func e(): ETTNd
17   func e(): ETTNd
18 }
19
20 class ETTNd {
21   // Parent link and children
22   var p: ETTNd
23   // Version of the node (for root only)
24   var v: In
25   // For storing non-spanning adjacent edges
26   val e = CTreeNode.S(e)
27   // Flags whether there are spanning
28   // or non-spanning edges in the subtree
29   var s: Bb
30   var n: Bb
31   // Concurrent removal operation (for root only)
32   var r: Synchronized
33   val l = L()
34 }
35
36 class RTreeNode      val u: ETTNd      val v: ETTNd {
37   var p: ETTNd? =      null
38 }
39
40 class SNode      val s: SNode      val e: In

```

**Listing 5:** Data structures used by our algorithm. Some fields, methods, and classes are omitted, because they are not important for the concurrent case.

**Lemma C.1.** *If `has_non_spanning_edges` flag management is performed according to Listing 6, where non-blocking additions just set needed flags to `true` and flag recalculation re-checks if it has written `false`, a thread holding the lock can not wrongly see `false` when there is a completely added non-spanning edge in the subtree.*

*Proof.* We prove the lemma by contradiction. Suppose that at some moment of time a situation, when a node has the flag set to `has` while there is an edge in its subtree, happened. Consider the first such moment. The situation can not be caused by a thread adding a non-spanning edge, because before the linearization point of the addition all flags in the path between a node and the root are set to `has`. Consequently, it was caused by a replacement search that recalculated the flag and set it to `has`. However, the recalculation should have re-checked the flag. Since all flag updates are performed from bottom to up the children of the node should have had correct values, so the re-check should have recognized the conflict and set the flag to `has`, which leads to a contradiction.  $\square$

**The `remove_edge(u, v)` operation.** The edge removal do not change much in comparison to the simplified case (Listing 7). The operation first reads the edge state (line 3), then checks whether the edge is already absent and it can just complete (line 5). If the edge is spanning or is about to become spanning, it is removed under the locks (lines 15-25). Otherwise, the operation tries to remove an edge without taking locks (line 10) by performing a CAS (line 32) and if

```

1 func Cññññññññ Eñ {
2   // Put the edge into the multisets and
3   // set all flags in the paths between
4   // the nodes and roots to 'true'.
5   this.ññññ.ññññ
6   ññññ this.ññññ)
7   this.ññññ.ññññ
8   ññññ this.ññññ)
9 }
10
11 func Cññññññññ Eñ {
12   // Remove a copy of an edge from the multisets
13   // without updating any flags.
14   this.ññññ.ññññ
15   this.ññññ.ññññ
16 }
17
18 func ññññ ETTNñ {
19   // Check if the flag is already enabled
20   if ññññ return
21   ññññ= true
22   ññññ= ññññ
23   // If is not a root, go up
24   if ññññ!= null: ññññ
25 }
26
27 func ññññ ETTNñ {
28   // Called only under locks.
29   // Calculate, whether there is a non-spanning edge
30   // in the subtree.
31   ññññ= ññññ
32   ññññ= ññññ
33   if (!ññññ:
34     // check whether 'true' was overridden by 'false'
35     ññññ= ññññ
36     if (ññññ:
37       | ññññ= true
38   }
39
40 func ññññ ETTNñ: Bññ=
41   ññññ> 0 or
42   ññññ or
43   ññññ

```

Listing 6: Edge information managing methods.  $\mathfrak{M}$  is a concurrent multiset.

successful, removes the edge physically (line 34). The replacement search upon spanning edge removal is discussed in a later paragraph.

```

1 func  $\text{add\_edge}(u, v)$  {
2   while (true):
3      $e := (u, v)$ 
4     // Check whether the edge is absent
5     if  $e == \text{null}$  or  $e == \text{INITIAL}$ : return
6      $e := \text{root}(u, v)$ 
7     if  $e == \text{SPANNING}$  or  $e == \text{IN\_PROGRESS}$ :
8       return  $e$ ,
9     if  $e == \text{NON-SPANNING}$ :
10      if  $\text{root}(u, v) == e$ :
11        return
12 }
13
14 func  $\text{remove\_edge}(u, v)$  {
15    $\text{root}(u, v)$  {
16      $e := (u, v)$ 
17     if  $e == \text{null}$  or  $e == \text{INITIAL}$ :
18       return // already absent
19     if  $e == \text{NON-SPANNING}$ :
20       return  $e$ ,
21     // Iterate over levels, search for a replacement,
22     // raise spanning and non-spanning edges
23      $\text{root}(u, v)$ 
24      $\text{remove\_edge}(u, v)$  // physical removal
25   }
26 }
27
28 func  $\text{remove\_edge}(u, v)$ 
29    $\text{root}(u, v)$  {
30   } :  $\text{root}(u, v)$  // returns whether is successful
31    $l := \text{root}(u, v)$ 
32   if  $l == \text{CAS}(\text{root}(u, v), \text{REMOVED}, l)$ :
33     // Finalize the removal.
34      $\text{root}(u, v)$ 
35     return true
36   return false
37 }

```

**Listing 7:** Implementation of the  $\text{remove\_edge}(u, v)$  operation.  $\text{root}(u, v)$  tries to remove a non-spanning edge without taking locks and can fail.

**The  $\text{add\_edge}(u, v)$  operation.** The basic algorithm of edge addition remains almost the same (Listing 8). The only differences are the fact that we use states instead of just statuses, the states are stored in a hash map, and if an operation sees that the edge is being added by a concurrent thread as spanning, it waits for the operation completion by taking the locks (lines 14-15). Besides, upon spanning edge addition the status changes twice – in the beginning and in the end of the operation to `IN_PROGRESS` and `SPANNING` respectively (lines 31 and 33).

The non-blocking addition (Listing 9) first optimistically adds information about the edge to corresponding multi-sets (line 5), then finds the root of the Cartesian tree (line 6), reads the `root` field from it that could have been published by a concurrent spanning edge removal in this component of connectivity. The ordering is important – similarly to Theorem 4.1, if the information is inserted before the removal, it will see the edge, so in the only interesting case the removal holds the locks while the addition is in lines 6-14 and the found root is correct. If the edge can be a replacement, it is proposed to a concurrent removal and the operation completes (lines 11-14). Otherwise, if the edge endpoints are still in the same component of connectivity, the operation is finalized by a status CAS (line 23).

An interesting change occurred in the replacement proposal. Previously, in the simplified case it was just a CAS to the replacement slot, but in the general case the logic is more difficult (lines 29-51). It reads the replacement slot (line 34) and then tries to update it in a lock-free style. If the removal is already over, nothing can be done (line 36). When the slot is empty, the operation just tries to write the edge via CAS (line 39). However, if there is another edge in the slot, the method can not just complete, since the replacement can be removed by a parallel thread before its status is changed

```

1 func addEdge {
2   // Create a new initial state to avoid the ABA problem.
3   h := StateINITIAL, 0)
4   // Atomically put the state only if there was none
5   // or get the present state.
6   h := compareAndSet h
7   if h == INITIAL:
8     h = h // help to add the edge
9   else if h != null: return // already present
10  while (true):
11    h := h
12    if h != h
13      // Synchronize with another thread if needed
14      if h == INPROGRESS:
15        lock {
16          return // someone inserted the edge
17        }
18        if h == h:
19          return
20      else:
21        lock
22        return
23  }
24
25 func addEdge h State {
26  lock {
27    // Check whether the edge is already present
28    if h != h return
29    if isSpanning != isSpanning:
30      // is a spanning edge
31      h = StateINPROGRESS, 0)
32      isSpanning
33      h = StateSPANNING, 0)
34    else:
35      // is a non-spanning edge.
36      // Simplified code of the non-blocking addition:
37      // Publish the edge before its addition
38      isSpanning
39      h = StateNON-SPANNING, 0)
40      if !h CAS(h h):
41        // already was non-spanning
42        isSpanning
43  }
44 }

```

Listing 8: Basic parts of the `addEdge`

operation. `addEdge` tries to add a non-spanning edge without taking locks and can fail.

tries to add a non-spanning

to SPANNING in line 13. So, in order not to wait for the status update, the thread can help by performing the same CAS (line 45). Because of such helping to avoid the ABA problem, the replacement slot contains not only the slot, but also its initial state. After that, in case of helping failure, the replacement slot is just cleared (line 50). This situation may seem bug-prone, because the removed edge can be non-spanning and still appropriate for the replacement. However, the fact that some thread managed to add the edge as non-spanning means that it started the addition and added the edge information before the removal, and thus, the removal will be able to find the information about this edge.

```

1 func ȳ
2 ȳ Eȳ ȳ Sȳ
3 ): Bȳ{ // returns whether is successful
4 // Publish the edge info before its addition.
5 Fȳ]ȳȳȳ
6 ȳ:= Fȳ]ȳȳȳ
7 ȳ:= ȳȳȳ
8 // Check whether there is a concurrent removal.
9 if ȳȳ!= null:
10 | if ȳȳȳȳ:
11 | | if ȳȳȳ ȳ ȳ:
12 | | | Fȳ]ȳȳȳ
13 | | | ȳ CAS(ȳ ȳ SȳSPANNING, 0))
14 | | | return true
15 | | else if ȳȳȳ== CLOSED:
16 | | | Fȳ]ȳȳȳ
17 | | | // The component is about to split,
18 | | | // so the edge will become spanning
19 | | | ȳȳȳ
20 | | | return true
21 | // Should the edge still be non-spanning?
22 | if Fȳ]ȳȳȳȳ ȳȳ and
23 | | ȳ CAS(ȳ ȳ SȳNON-SPANNING, 0)):
24 | | | return true
25 | ȳȳȳ
26 | return false
27 }
28
29 func ȳ
30 ȳ Sȳȳȳȳȳȳȳ
31 ȳ Eȳ ȳ Sȳ
32 ): Bȳ{
33 while (true):
34 | ȳ:= ȳȳȳ
35 | // Is the replacement search already completed?
36 | if ȳ== CLOSED: return false
37 | if ȳ== null:
38 | | // Try to put the edge in the slot.
39 | | if ȳȳ CAS(null, (ȳ ȳ):
40 | | | return true
41 | | | continue // try again
42 | if ȳȳ== ȳ
43 | | return true // already is the replacement
44 | | (ȳ ȳ := ȳ
45 | | if ȳ CAS(ȳ ȳ SȳSPANNING, 0))
46 | | or ȳȳ== SPANNING:
47 | | | return false
48 | | // Replacement edge was removed or is NON-SPANNING,
49 | | // so can clear the slot.
50 | | ȳȳ CAS(ȳ ȳ null)
51 }

```

**Listing 9:** Non-blocking addition of a non-spanning edge.

**Replacement Search.** Given all the ideas we presented before, the replacement search is rather straightforward. It iterates over all current-level non-spanning edges using `ȳ` flags and the concurrent



multi-sets. For each edge it checks the status and the level of the edge. Upon finding an edge in the `INITIAL` status, i.e. an edge that is being added by a parallel thread as non-spanning it helps to add it (lines 13-27). After that, the search checks whether the edge can be a replacement and in this case writes it to its own replacement slot via the same `add` method (line 31) and, in case of success, deletes the information about the edge. If the edge is not the replacement, to amortize the cost of its finding the edge is promoted to the next ETT level (lines 37-43). This promotion is optimistic – it adds the information about the edge to the next level before the CAS of its state and then removes the edge information from either the current level or the next level depending on the CAS results. When the removal finishes its search, it “closes” the replacement slot via the `close` method (line 53).

**Linearization Points.** The correctness of our algorithm was discussed in Theorem 4.1 in the simplified case and the arguments remain the same for the general case. We verified our algorithms with stress testing and bounded model checking. Since our algorithm is complex, we would also like to list all possible linearization points for every *successful* operation. For blocking edge modifications, we consider only linearization points of connected component changes.

- `add`  $\psi$ 
  - spanning edge
    - \* The same linearization point as for the ETT addition (Figure 2).
    - \* The successful CAS in line 13 of Listing 9.
    - \* The successful CAS by a concurrent thread in line 45 of Listing 9.
    - \* The successful CAS by a concurrent thread in line 18 of Listing 10.
  - non-spanning edge
    - \* The successful state CAS in line 23 of Listing 9.
    - \* The successful state CAS in line 40 of Listing 8.
    - \* The successful state CAS by a concurrent thread in line 24 of Listing 10.
- `remove`  $\psi$ 
  - spanning edge
    - \* If there is no replacement in  $F_0$ , then the linearization point is the same as for the ETT removal (Figure 3), otherwise components of connectivity do not change.
  - non-spanning edge
    - \* The successful state CAS in line 32 of Listing 7.
- `close`  $\psi$ 
  - See Theorem 3.2, where we proved the existence of a linearization point.

```

1 // Replacement search for F_0
2 func 重
3  重 ETTN 重 R 重
4 ): B 重{
5   if 重== null or ! 重
6     return false
7   重:= false
8   // Iterate over all non-spanning edges.
9   for 重 in 重
10    重:= 重
11    if 重== null: continue
12    if 重!= 0: continue
13    if 重== INITIAL:
14      // Help to add the edge.
15      if 重:
16        // Add as a replacement.
17        if 重 重:
18          if 重 CAS(重 重 SPANNING, 0):
19            重= true, break
20        // Should the edge still be non-spanning?
21        if FQ] 重 == FQ] 重:
22          FQ] 重
23          重 = SNON-SPANNING, 0)
24          if 重 CAS(重 重 重:
25            重= 重 // successfully added
26          else:
27            FQ] 重 // failure
28        if 重!= NON-SPANNING: continue
29        if 重:
30          if 重 CAS(重 重 SPANNING, 0):
31            if 重 重:
32              FQ] 重 // replacement proposed
33            else:
34              重= 重 // revert the state
35            重= true, break
36        else:
37          // Promote the non-spanning edge to the next level
38          FQ] 重
39          重:= SNON-SPANNING, 1)
40          if 重 CAS(重 重 重:
41            FQ] 重 // promotion successful
42          else
43            FQ] 重 // promotion failed
44        if ! 重
45          重= 重 重 重
46        if ! 重
47          重= 重 重 重
48        重
49        return 重
50 }
51
52 // Closes the slot and returns a replacement or null
53 func 重
54 重 R 重
55 ): E 重{
56   if 重 CLOSED):
57     return null // the slot is closed
58   return 重 重
59 }

```

Listing 10: Replacement search for  $F_0$ .