

Parallel Batch-Dynamic Graph Connectivity

Umut A. Acar
Carnegie Mellon University
umut@cs.cmu.edu

Daniel Anderson
Carnegie Mellon University
dlanders@cs.cmu.edu

Guy E. Blelloch
Carnegie Mellon University
guyb@cs.cmu.edu

Laxman Dhulipala
Carnegie Mellon University
ldhulipa@cs.cmu.edu

ABSTRACT

In this paper, we study batch parallel algorithms for the dynamic connectivity problem, a fundamental problem that has received considerable attention in the sequential setting. The best sequential algorithm for dynamic connectivity is the elegant level-set algorithm of Holm, de Lichtenberg and Thorup (HDT), which achieves $O(\lg^2 n)$ amortized time per edge insertion or deletion, and $O(\lg n)$ time per query.

We design a parallel batch-dynamic connectivity algorithm that is work-efficient with respect to the HDT algorithm for small batch sizes, and is asymptotically faster when the average batch size is sufficiently large. Given a sequence of batched updates, where Δ is the average batch size of all deletions, our algorithm achieves $O(\lg n \lg(1 + n/\Delta))$ expected amortized work per edge insertion and deletion and $O(\lg^3 n)$ depth w.h.p. Our algorithm answers a batch of k connectivity queries in $O(k \lg(1 + n/k))$ expected work and $O(\lg n)$ depth w.h.p. To the best of our knowledge, our algorithm is the first parallel batch-dynamic algorithm for connectivity.

ACM Reference Format:

Umut A. Acar, Daniel Anderson, Guy E. Blelloch, and Laxman Dhulipala. 2019. Parallel Batch-Dynamic Graph Connectivity. In *31st ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '19)*, June 22–24, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3323165.3323196>

1 INTRODUCTION

Computing the connected components of a graph is a fundamental problem that has been studied in many different models of computation [6, 8, 33, 53, 56, 62]. The **connectivity problem** takes as input an undirected graph G and requires an assignment of labels to vertices such that two vertices have the same label if and only if they are in the same connected component. The dynamic version of the problem requires maintaining a data structure over an n vertex undirected graph that supports insertions and deletions of edges, and queries of whether two vertices are in the same connected component. Despite the large body of work on the dynamic connectivity problem over the past two decades [22, 29, 30, 33, 34, 38, 40, 45, 63, 64, 69, 70], little is known about batch-dynamic connectivity algorithms that process *batches of queries and updates*, either sequentially or in parallel.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SPAA '19, June 22–24, 2019, Phoenix, AZ, USA
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6184-2/19/06...\$15.00
<https://doi.org/10.1145/3323165.3323196>

Traditional dynamic algorithms were motivated by applications where data undergoes small changes that can be adequately handled by updates of single elements. Today, however, applications operate on increasingly large datasets that undergo rapid changes over time: for example, millions of individuals can simultaneously interact with a web site, make phone calls, send emails and so on. In the context of these applications, traditional dynamic algorithms require serializing the changes made and processing them one at a time, missing an opportunity to exploit the parallelism afforded by processing batches of changes.

Motivated by such applications, there has been recent interest in developing theoretically efficient parallel batch-dynamic algorithms [1, 3, 59, 65]. In the batch-dynamic setting, instead of applying one update or query at a time, a whole batch is applied. A batch could be of size $\lg n$, \sqrt{n} , or $n/\lg n$ for example. There are two advantages of applying operations in batches.

- (1) Batching operations allows for more parallelism.
- (2) Batching operations can reduce the cost of each update.

In this paper we are interested in both these advantages. We use the term parallel batch-dynamic to mean algorithms that process batches of operations instead of single ones, and for which the algorithm itself is parallel. The underlying parallel model used in this paper is a formalization of the widely used shared-memory work-depth model [9, 11, 14, 17].

Understanding the connectivity structure of graphs is of significant practical interest, for example, due to its use as a primitive for clustering the vertices of a graph [54]. Due to the importance of connectivity there are several implementations of parallel batch-dynamic connectivity algorithms [35, 36, 43, 55, 67, 68]. In the worst case, however, these algorithms may recompute the connected components of the entire graph even for very small batches. Since this requires $O(m+n)$ work, it makes the worst-case performance of the algorithms no better than running a static parallel algorithm. On the theoretical side, existing batch-dynamic efficient connectivity algorithms have only been designed for restricted settings, e.g., in the incremental setting when all updates are edge insertions [59], or when the underlying graph is a forest [1, 52, 65].

Therefore, two important questions are:

- (1) *Is there a batch-dynamic connectivity algorithm that is asymptotically faster than existing dynamic connectivity algorithms for large enough batches of insertions, deletions and queries?*
- (2) *Can the batch-dynamic connectivity algorithm be parallelized to achieve low worst-case depth?*

In this paper we give an algorithm that answers both of these questions affirmatively. To simplify our exposition and present the main ideas, we first give a less efficient version of the algorithm

that runs in $O(\lg^4 n)$ depth w.h.p.¹ and performs $O(\lg^2 n)$ expected amortized work per update, making it work-efficient with respect to the sequential algorithm of Holm, de Lichtenberg, and Thorup. Next, we describe the improved algorithm which runs in $O(\lg^3 n)$ depth w.h.p. and achieves an improved work bound that is asymptotically faster than the HDT algorithm for sufficiently large batch sizes. We note that our depth bounds hold even when processing the updates one at a time, ignoring batching. Our improved work bounds are derived by a novel analysis of the work performed by the algorithm over all batches of deletions.

Our contribution is summarized by the following theorem:

THEOREM 1.1. *There is a parallel batch-dynamic data structure which, given batches of edge insertions, deletions, and connectivity queries processes all updates in $O\left(\lg n \lg\left(1 + \frac{n}{\Delta}\right)\right)$ expected amortized work per edge insertion or deletion where Δ is the average batch size of a deletion operation. The cost of connectivity queries is $O(k \lg(1 + n/k))$ work and $O(\lg n)$ depth for a batch of k queries. The depth to process a batch of edge insertions and deletions is $O(\lg n)$ and $O(\lg^3 n)$ respectively.*

Technical Overview. The starting point of our algorithm is the classic Holm, de Lichtenberg and Thorup (HDT) dynamic connectivity algorithm [33]². Like nearly all existing dynamic connectivity algorithms, the HDT algorithm maintains a spanning forest certifying the connectivity of the graph. The algorithm maintains a set of $\lg n$ nested forests under two carefully designed invariants. The forests are represented using the Euler tour tree (ET-tree) data structure [29, 44].

The main challenge in a dynamic connectivity algorithm is to efficiently find a *replacement edge*, or a non-tree edge going between the two disconnected components after deleting a tree edge. The key idea of the HDT algorithm is to organize the spanning-forest of the graph into $\lg n$ levels of trees. The top-most level of the structure stores a spanning forest of the entire graph, and each level contains all tree-edges stored in levels below it. The algorithm ensures that the largest size of a component at level i is 2^i . Using these invariants, the algorithm is able to cleverly search the tree edges so that each non-tree edge is examined at most $\lg n$ times as a candidate replacement edge. The main idea is to store each non-tree edge at a single level (initially the top-most level), and push the edge to a lower level each time it is unsuccessfully considered as a replacement edge. Since there are $\lg n$ levels, and the cost of discovering, processing, and removing an edge from each level using ET-tree operations is $O(\lg n)$, the amortized cost of the HDT algorithm is $O(\lg^2 n)$ per edge operation. We now discuss the main challenges and sequential bottlenecks that arise in the HDT algorithm that a parallel batch-dynamic algorithm must address.

Efficiently searching for replacements: A challenge, and sequential bottleneck in the HDT algorithm is the fact that it processes each non-tree edge one at a time—a property which is crucial for achieving good amortized bounds. Aside from hindering parallelism, processing the edges one at a time eliminates any potential for improved batch bounds, since finding the representative of the

endpoints of an edge costs $O(\lg n)$ time per query. Therefore, to obtain an efficient batch or parallel algorithm we must examine batches of *multiple* non-tree edges at a time, while also ensuring that we do not perform extra work that cannot be charged to level-decreases on an edge. Our approach is to use a *doubling technique*, where we examine sets of non-tree edges with geometrically increasing sizes.

Handling Batches of Deletions: Another challenge is that processing a batch of deletions can shatter a component into multiple disconnected pieces. Since the HDT algorithm deletes at most a single tree edge per deletion operation, it handles exactly two disconnected pieces per level. In contrast, since we delete batches of edges in our batch-dynamic algorithm, we may have many disconnected pieces at a given level, and must search for replacement edges reconnecting these pieces. Our algorithm searches for a replacement edge from each piece that is small enough to be pushed down to the next lower level.

Each round of both of our algorithms can be viewed as calling an oracle which returns a set of replacement edges incident on the disconnected pieces that we are trying to reconnect. Unlike in the HDT algorithm which terminates once it finds any replacement edge, the edges returned by the oracle may not fully restore the connectivity of the original component. In particular, the replacement edges that we find may contain multiple edges going between the same pieces (like in Borůvka’s algorithm) or may contain cycles, which must be dealt with since each level of the data structure represents a forest. Our approach to handling these issues is to run a static spanning forest algorithm on the replacement edges found in this round, and insert only the spanning forest edges into the ET-tree at the current level.

Both of our algorithms alternate between a first phase which calls the oracle to find a set of replacement edges, and a second phase which determines a set of replacement edges that can be committed as tree edges. The difference is that our first (simpler) algorithm (Section 3) requires $O(\lg^2 n)$ oracle queries per level, whereas the second algorithm (Section 4) only requires $O(\lg n)$ oracle queries due to a more careful doubling technique.

Dynamic trees supporting batching: Another obstacle to improving on the bounds of the HDT algorithm is that the classic ET-tree data structure performs links and cuts one at a time. To achieve good batch bounds for forest operations, we use a recently developed solution to the batch-parallel forest connectivity problem by Tseng et al. [65]. Their data structure, which we refer to as a batch-parallel ET-tree processes a set of k links, cuts, or connectivity queries in $O(k \lg(1 + n/k))$ work and $O(\lg n)$ depth. We extend the data structure to supports operations such as fetching the first l non-tree edges in the tree efficiently.

2 PRELIMINARIES

Model. In this paper we analyze our algorithms in the work-depth model using fork-join style parallelism. Specifically, we use a particular work-depth model called the *Threaded Random Access Machine* (TRAM), which is closely related to the PRAM but more closely models current machines and programming paradigms that are asynchronous and support dynamic forking. The model can work-efficiently cross-simulate a CRCW PRAM, equipped with the same

¹ We say that an algorithm has $O(f(n))$ cost **with high probability (w.h.p.)** if it has $O(k \cdot f(n))$ cost with probability at least $1 - 1/n^k$, $k \geq 1$.

²We provide full details of the HDT algorithm in Section 2.2.

atomic instructions, and is therefore essentially equivalent to the classic CRCW PRAM model. We formally define the model and provide more details about the simulations in the full version of our paper [2], and refer the interested reader to [12] for full details.

Our algorithms are designed using *nested fork-join parallelism* in which a procedure can *fork* off another procedure call to run in parallel and then wait for forked calls to complete with a *join* synchronization [11]. Our efficiency bounds are stated in terms of work and depth, where **work** is the total number of vertices in the thread DAG and where **depth (span)** is the length of the longest path in the DAG [11].

Parallel Primitives. The following parallel procedures are used throughout the paper. A **semisort** takes an input array of elements, where each element has an associated key and reorders the elements so that elements with equal keys are contiguous, but elements with different keys are not necessarily ordered. The purpose is to collect equal keys together, rather than sort them. Semisorting a sequence of length n can be performed in $O(n)$ expected work and $O(\lg n)$ depth w.h.p. assuming access to a uniformly random hash function mapping keys to integers in the range $[1, n^{O(1)}]$ [26, 51].

A **parallel dictionary** data structure supports batch insertion, batch deletion, and batch lookups of elements from some universe with hashing. Gil et al. describe a parallel dictionary that uses linear space and achieves $O(k)$ work and $O(\lg^* k)$ depth w.h.p. for a batch of k operations [25].

The **pack** operation takes an n -length sequence A and an n -length sequence B of booleans as input. The output is a sequence A' of all the elements $a \in A$ such that the corresponding entry in B is *true*. The elements of A' appear in the same order that they appear in A . Packing can be easily implemented in $O(n)$ work and $O(\lg n)$ depth [37].

Useful Lemmas. The following lemmas are useful for analyzing the work bounds of our parallel algorithms. We provide proofs in the full version of our paper [2].

LEMMA 2.1. *Let n_1, n_2, \dots, n_c and k_1, k_2, \dots, k_c be sequences of non-negative integers such that $\sum k_i = k$, and $\sum n_i = n$. Then*

$$\sum_{i=1}^c k_i \lg \left(1 + \frac{n_i}{k_i} \right) \leq k \lg \left(1 + \frac{n}{k} \right).$$

LEMMA 2.2. *For any non-negative integers n and r ,*

$$\sum_{w=0}^r 2^w \lg \left(1 + \frac{n}{2^w} \right) = O \left(2^r \lg \left(1 + \frac{n}{2^r} \right) \right).$$

LEMMA 2.3. *For any $n \geq 1$, the function $x \lg \left(1 + \frac{n}{x} \right)$ is strictly increasing with respect to x for $x \geq 1$.*

2.1 Batch-Dynamic Trees

The batch-dynamic trees problem is to represent a forest as it undergoes batches of links, cuts, and connectivity queries. A **link** operation inserts an edge connecting two trees in the forest. A **cut** deletes an edge from the forest, breaking one tree into two trees. A **connectivity** query returns whether two vertices are connected by a path (in the same tree) in the forest. The interface is formally defined as follows:

Batch-Dynamic Trees Interface.

- **BATCHLINK**($\{(u_1, v_1), \dots, (u_k, v_k)\}$) takes a sequence of edges and adds them to the graph G . The input edges must not create a cycle in G .
- **BATCHCUT**($\{(u_1, v_1), \dots, (u_k, v_k)\}$) takes a sequence of edges and removes them from the graph G .
- **BATCHCONNECTED**($\{(u_1, v_1), \dots, (u_k, v_k)\}$) takes a sequence of tuples representing queries. The output is a sequence where the i -th entry returns whether vertices u_i and v_i are connected by a path in G .
- **BATCHFINDREPR**($\{(x_1, \dots, x_k)\}$) takes a sequence of pointers to tree elements. The output is a sequence where the i -th entry is the *representative* (repr) of the tree in which x_i lives. The representative is defined so that $\text{repr}(u) = \text{repr}(v)$ if and only if u and v are in the same tree. Note that representatives are invalidated after the tree is modified.

Batch-Parallel Euler Tour Trees. In this paper we make use of a recently developed, parallel solution to the batch-dynamic trees problem, called a batch-parallel Euler tour tree (batch-parallel ET-trees) [65]. The data structure represents each ET-tree sequence using a concurrent skip-list, and reduces bulk link, cut, and query operations to bulk operations on the concurrent skip-list.

Tseng et al. [65] prove the following theorem on the efficiency of the batch-parallel ET-tree:

THEOREM 2.4. *A batch of k links, k cuts, k connectivity queries, or k representative queries over an n -vertex forest can be processed in $O(k \lg(1 + n/k))$ expected work and $O(\lg n)$ depth with high probability.*

The trees also support augmentation with an associative and commutative function $f : D^2 \rightarrow D$ with values from D assigned to vertices and edges of the forest. The goal is to compute f over subtrees of the represented forest. The interface can be easily extended with the following batch-parallel primitives for updating and querying augmented values.

The full version of our paper [2] contains information about additional tree operations that are needed to efficiently implement our algorithms.

2.2 The sequential (HDT) algorithm

Our parallel algorithm is based on the sequential algorithm of Holm, De Lichtenberg, and Thorup [33], which we refer to as the HDT algorithm. The HDT algorithm assigns to each edge in the graph, an integer *level* from 1 to $\lg n$. The levels correspond to sequence of subgraphs $G_1 \subset G_2 \subset \dots \subset G_{\lg n} = G$, such that G_i contains all edges with level at most i . The algorithm also maintains a spanning forest F_i of each G_i such that $F_1 \subset F_2 \subset \dots \subset F_{\lg n}$. Each forest is maintained using a set of augmented ET-trees which we describe shortly. Throughout the algorithm, the following invariants are maintained.

INVARIANT 1. $\forall i = 1 \dots \lg n$, the connected components of G_i have size at most 2^i .

INVARIANT 2. $F_{\lg n}$ is a minimum spanning forest where the weight of each edge is its level.

Connectivity Queries. To perform a connectivity query in G , it suffices to query $F_{\lg n}$, which takes $O(\lg n)$ time by querying for the root of each Euler tour tree and returning whether the roots are equal. We note that in [33], a query time of $O(\lg n / \lg \lg n)$ is achieved by storing the Euler tour of $F_{\lg n}$ in a B-tree with branching factor $\lg n$.

Inserting an Edge. An edge insertion is handled by assigning the edge to level $\lg n$. If the edge connects two currently disconnected components, then it is added to $F_{\lg n}$.

Deleting an Edge. Deletion is the most interesting part of the algorithm. If the deleted edge is not in the spanning forest $F_{\lg n}$, the algorithm removes the edge and does nothing to $F_{\lg n}$ as the connectivity structure of the graph is unchanged. Otherwise, the component containing the edge is split into two. The goal is to find a **replacement edge**, that is, an edge crossing the split component.

If the deleted edge had level i , then the *smaller* of the two resulting components is searched starting at level i in order to locate a replacement edge. Before searching this component, all tree edges whose level is equal to i have their level decremented by one. As the smaller of the split components at level i has size $\leq 2^{i-1}$, pushing the entire component to level $i - 1$ does not violate Invariant 1. Next, the non-tree edges at level i are considered one at a time as possible replacement edges. Each time the algorithm examines an edge that is not a replacement edge, it decreases the level of the edge by one. If no replacement is found, it moves up to the next level and repeats. Note that because the algorithm first pushes all tree edges to level $i - 1$, any subsequent non-tree edges that may be pushed from level i to level $i - 1$ will not violate Invariant 2.

Implementation and Cost. To efficiently search for replacement edges, the ET-trees are augmented with two additional pieces of information. The first augmentation is to maintain the number of non-tree edges whose level equals the level of the tree. The second augmentation maintains the number of tree-edges whose level is equal to the level of the tree.

Using these augmentations, each successive non-tree edge (or tree edge) whose level is equal to the level of the tree can be found in $O(\lg n)$ time. Furthermore, checking whether the edge is a replacement edge can be done in $O(\lg n)$ time. Lastly, the cost of pushing an edge that is not a replacement edge to the lower level is $O(\lg n)$, since it corresponds to inserting the edge into an adjacency structure and updating the augmented values. Since each edge can be processed at most once per level, paying a cost of $O(\lg n)$, and there are $\lg n$ levels, the overall amortized cost per edge is $O(\lg^2 n)$.

3 A PARALLEL ALGORITHM

In this section, we give a simple parallel batch-dynamic connectivity algorithm based on the HDT algorithm. The underlying invariants maintained by our parallel algorithm are identical to the sequential HDT algorithm: we maintain $\lg n$ levels of spanning forests subject to Invariants 1 and 2. The main challenge, and where our algorithm departs from the HDT algorithm is in how we search for replacement edges in parallel, and how we search multiple components in parallel. We show by a charging argument that this parallel algorithm is work-efficient with respect to the HDT algorithm—it performs $O(\lg^2 n)$ amortized work per edge insertion or deletion. Furthermore, we show that the depth of this algorithm is $O(\lg^4 n)$.

Although these bounds are subsumed by the improved parallel algorithm we describe in Section 4, the parallel algorithm in this section is useful to illustrate the main ideas in this paper.

Data Structures. Each spanning forest, F_i , is represented using a set of batch-parallel ET-trees [65]. We represent the edges of the graph in a parallel dictionary E_D for convenience (see Section 2). We also store an adjacency array, $A_i[u]$, at each level i , and for each vertex u to store the tree and non-tree edges incident on u with level i . Note that tree and non-tree edges are stored separately so that they can be accessed separately. The adjacency arrays support batch insertion and deletion of edges, as well as the ability to fetch a batch of edges of a desired size. These operations have the following cost bounds.

LEMMA 3.1. *INSERTEDGES, DELETEDGES, and FETCHEDGES can be implemented in $O(1)$ amortized work per edge and in $O(\lg n)$ depth.*

We refer the reader to the full version of our paper [2] for proofs and full details on the adjacency data structure.

3.1 Connectivity Queries

As in the sequential algorithm, a connectivity query can be answered by simply performing a query on $F_{\lg n}$. Algorithm 1 gives pseudocode for the batch connectivity algorithm. The bound we achieve follows from the batch bounds on batch-parallel ET-trees.

Algorithm 1 The batch query algorithm

```

1: procedure BATCHQUERY( $\{(u_1, v_1), (u_2, v_2), \dots, (u_k, v_k)\}$ )
2:   return  $F_{\lg n}$ .BATCHQUERY( $\{(u_1, v_1), (u_2, v_2), \dots, (u_k, v_k)\}$ )

```

THEOREM 3.2. *A batch of k connectivity queries can be processed in $O\left(k \lg\left(1 + \frac{n}{k}\right)\right)$ expected work and $O(\lg n)$ depth w.h.p.*

PROOF. Follows from Theorem 2.4. □

3.2 Inserting Batches of Edges

To perform a batch insertion, we first determine a set of edges in the batch that increase the connectivity of the graph. To do so, we treat each current connected component of the graph as a vertex, and build a spanning forest of the edges being inserted over this contracted graph. The edges in the resulting spanning forest are then inserted into the topmost level in parallel.

Algorithm 2 The batch insertion algorithm

```

1: procedure BATCHINSERT( $U = \{(u_1, v_1), \dots, (u_k, v_k)\}$ )
2:   For all  $e_i \in U$ , set  $l(e_i) \leftarrow \lg n$  in parallel
3:   Update  $A_{\lg n}[u]$  for edges incident on  $u$ 
4:    $R \leftarrow \{(F_{\lg n}.\text{FINDREPR}(u), F_{\lg n}.\text{FINDREPR}(u)) \mid (u, v) \in U\}$ 
5:    $T' \leftarrow \text{SPANNINGFOREST}(R)$ 
6:    $T \leftarrow$  edges in  $U$  corresponding to  $T'$ 
7:   Promote edges in  $T$  to tree edges
8:    $F_{\lg n}$ .BATCHINSERT( $T$ )

```

Algorithm 2 gives pseudocode for the batch insertion algorithm. We assume that the edges given as input in U are not present in the graph. Each vertex u that receives an updated edge inserts its edges into $A_{\lg n}[u]$ (Line 3). This step can be implemented by first running a semisort to collect all edges incident on u .

The last step is to insert edges that increase the connectivity of the graph as tree edges (Lines 4–8). The algorithm starts by computing the representatives for each edge (Line 4). The output is an array of edges, R , which maps each original (u, v) edge in U to $(\text{FINDREPR}(u), \text{FINDREPR}(v))$ (note that these calls can be batched using BATCHFINDREPR). Next, it computes a spanning forest over the tree edges (Line 5). Finally, the algorithm promotes the corresponding edges in U to tree edges. This is done by updating the appropriate adjacency lists and inserting them into $F_{\lg n}$ (Lines 7–8).

THEOREM 3.3. *A batch of k edge insertions can be processed in $O(k \lg(1 + \frac{n}{k}))$ expected work and $O(\lg n)$ depth w.h.p.*

PROOF. Lines 2–3 cost $O(k)$ work and $O(\lg k)$ depth w.h.p. using our bounds for updating A (see Lemma 3.1). The find representative queries (Line 4) can be implemented using a BATCHFINDREPR call, which costs $O(k \lg(1 + \frac{n}{k}))$ expected work and $O(\lg n)$ depth w.h.p. by Theorem 2.4. Computing a spanning forest (Line 5) can be done in $O(k)$ expected work and $O(\lg k)$ depth w.h.p. using Gazit’s connectivity algorithm [24]. Finally, updating the adjacency lists and inserting the spanning forest edges into $F_{\lg n}$ costs $O(k \lg(1 + \frac{n}{k}))$ expected work and $O(\lg n)$ depth w.h.p. (Lines 7–8). \square

3.3 Deleting Batches of Edges

As in the sequential HDT algorithm, searching for replacement edges after deleting a batch of tree edges is the most interesting part of our parallel algorithm. A natural idea for parallelizing the HDT algorithm is to simply scan all non-tree edges incident on each disconnected component in parallel. Although this approach has low depth per level, it may examine a huge number of candidate edges, but only push down a few non-replacement edges. In general, it is unable to amortize the work performed checking all candidates edges at a level to the edges that experience level decreases. To amortize the work properly while also searching the edges in parallel we must perform a more careful exploration of the non-tree edges. Our approach is to use a *doubling* technique, in which we geometrically increase the number of non-tree edges explored as long as we have not yet found a replacement edge. We show that using the doubling technique, the work performed (and number of non-tree edges explored) is dominated by the work of the last phase, when we either find a replacement edge, or run out of non-tree edges. Our amortized work-bounds follow by a per-edge charging argument, as in the analysis of the HDT algorithm.

The Deletion Algorithm. Algorithm 3 shows the pseudocode for our parallel batch deletion algorithm. As with the batch insertion algorithm, we assume that each edge is present in U in both directions. Given a batch of k edge deletions, the algorithm first deletes the given edges from their respective adjacency lists in parallel (Line 2). It then filters out the tree edges (Line 3) and deletes each tree edge e from $F_i \dots, F_{\lg n}$, where i is the level of e (Line 4). Next, it computes C , a set of *components* (representatives) from the deleted tree edges (Line 5). For each deleted tree edge, e , the algorithm includes the representatives of both endpoints in the forest at $l(e)$, which must be in different components as e is a deleted tree edge. Finally, the algorithm loops over the levels, starting at the

lowest level where a tree edge was deleted (Line 7), and calls $\text{PARALLELLEVELSEARCH}$ at each level. Each call to $\text{PARALLELLEVELSEARCH}$ takes i , the level to search, C , the current set of disconnected components, and S , an initially empty set of replacement edges that the algorithm discovers over the course of the searches (Line 8)

Algorithm 3 The batch deletion algorithm

```

1: procedure BATCHDELETION( $U = \{e_1, \dots, e_k\}$ )
2:   Delete  $e \in U$  from  $A_0, \dots, A_{\lg n}$ 
3:    $T \leftarrow \{e \in U \mid e \in F_{\lg n}\}$  ▷ Tree edges to delete
4:   Delete  $e \in T$  from  $F_0, \dots, F_{\lg n}$ 
5:    $C \leftarrow \cup_{e=(u,v) \in T} (F_{l(e)}. \text{FINDREPR}(u), F_{l(e)}. \text{FINDREPR}(v))$ 
6:    $S \leftarrow \emptyset$ 
7:   for  $i \in [\min_l \leftarrow \min_{e \in T, \lg n}]$  do
8:      $(C, S) \leftarrow \text{PARALLELLEVELSEARCH}(i, C, S)$ 

```

Algorithm 4 The parallel level search algorithm

```

1: procedure COMPONENTSEARCH( $i, c$ )
2:    $w \leftarrow 1, w_{\max} \leftarrow c. \text{NUMNONTREEEDGES}$ 
3:   while  $w < 2w_{\max}$  do
4:      $w \leftarrow \min(w, w_{\max})$ 
5:      $E_c \leftarrow$  First  $w$  non-tree edges in  $c$ 
6:     Push all non-replacement edges in  $E_c$  to level  $i - 1$ 
7:     if  $E_c$  contains a replacement edge then
8:       return  $\{r\}$ , where  $r$  is any replacement edge in  $E_c$ 
9:      $w \leftarrow 2w$ 
10:  return  $\emptyset$ 
11: procedure PARALLELLEVELSEARCH( $i, L = \{c_1, c_2, \dots\}, S$ )
12:   $F_i. \text{BATCHINSERT}(S)$ 
13:   $C \leftarrow c \in L$  with size  $\leq 2^{i-1}$ 
14:   $D \leftarrow c \in L$  with size  $> 2^{i-1}$ 
15:  while  $|C| > 0$  do
16:    Push level  $i$  tree edges of components in  $C$  to level  $i - 1$ 
17:     $R \leftarrow \cup_{c \in C} \text{COMPONENTSEARCH}(i, c)$  ▷ In parallel
18:     $R' \leftarrow \{(F_i. \text{FINDREPR}(u), F_i. \text{FINDREPR}(v)) \mid (u, v) \in R\}$ 
19:     $T' \leftarrow \text{SPANNINGFOREST}(R')$ 
20:     $T \leftarrow$  Edges in  $R$  corresponding to edges in  $T'$ 
21:    Promote edges in  $T$  to tree edges
22:     $F_i. \text{BATCHINSERT}(T)$ 
23:     $S \leftarrow S \cup T$ 
24:     $C \leftarrow \{F_i. \text{REPR}(c) \mid c \in C\}$ 
25:     $Q \leftarrow \{c \in C \text{ with no non-tree edges, or size } > 2^{i-1}\}$ 
26:     $D \leftarrow D \cup Q$ 
27:     $C \leftarrow C \setminus Q$ 
28:  return  $(D, S)$ 

```

Searching a Level in Parallel. The bulk of the work done by the deletion algorithm is performed by Algorithm 4, which implements a subroutine that searches the disconnected components at a given level of the data structure in parallel. The input to $\text{PARALLELLEVELSEARCH}$ is an integer i , the level to search, a set of representatives of the disconnected components, L , and the set of replacement spanning forest edges that were found in levels lower than i , S . The output of $\text{PARALLELLEVELSEARCH}$ is the set of components that are still disconnected after considering the non-tree edges at this level, and the set of replacement spanning forest edges found so far.

PARALLELEVESEARCH starts by inserting the new spanning forest edges in S into F_i (Line 12). Next, it computes C and D , which are the components that are active and inactive at this level, respectively (Lines 13–14). The main loop of the algorithm (Lines 15–27) operates in a number of **rounds**.

Each round first pushes down all tree edges at level i of every active component. It then finds a single replacement edge incident to each active component, searching the active components in parallel, pushing any non-replacement edge to level $i - 1$. It then promotes a maximal acyclic subset of the replacement edges found in this round to tree edges, and proceeds to the next round. The rounds terminate once all components at this level are deactivated by either becoming too large to search at this level, or because the algorithm finished examining all non-tree edges incident to the component at this level.

The main loop (Lines 15–27) works as follows. The algorithm first pushes any level i tree edges in an active component down to level $i - 1$. The active components in C have size at most 2^{i-1} , meaning that any tree edges they have at level i can be pushed to level $i - 1$ (Line 16) without violating Invariant 1. Next, the algorithm searches each active component for a replacement edge in parallel by calling the COMPONENTSEARCH procedure in parallel over all components (Line 17). This procedure either returns an empty set if there are no replacement edges incident to the component, or a set containing a single replacement edge. Next, the algorithm maps the replacement edge endpoints to their current component’s representatives by calling FINDREPR on each endpoint (Line 18). It then computes a spanning forest over these replacement edges (Line 19) and maps the edges included in the spanning forest back to their original endpoints ids (Line 20). Observe that the edges in T constitute a maximal acyclic subset of replacement edges of R in F_i . The algorithm therefore promotes the edges in T to tree edges (Lines 21–22). Note that the new tree edges are not immediately inserted into all higher level spanning trees. Instead, the edges are buffered by adding them to S (Line 23) so that they will be inserted when the higher level is reached in the search. Finally, the algorithm updates the set of components by computing their representatives on the updated F_i (Line 24), and filtering out any components which have no remaining non-tree edges, or become larger than 2^{i-1} (i.e., become unsearchable at this level) into D (Lines 30–27).

We now describe the COMPONENTSEARCH procedure (Lines 1–10). The search consists of a number of **phases**, where the i ’th phase searches the first 2^i non-tree edges, or all of the non-tree edges if 2^i is larger than the number of non-tree edges in c . The search terminates either once a replacement edge incident to c is found (Line 7), or once the algorithm unsuccessfully examines all non-tree edges incident to c (Line 3). Initially w , the search size, is set to 1 (Line 4). On each phase, the algorithm retrieves the first w many non-tree edges, E_c (Line 5). It pushes all non-tree edges that are not replacements to level $i - 1$ (Line 6). It then checks whether any of the edges in E_c are a replacement edge, and if so, returns one of the replacement edges in E_c (Line 8). Note that checking whether an edge is a replacement edge is done using BATCHFINDREPR. Otherwise, if no replacement edge was found it doubles w (Line 9) and continues.

Cost Bounds. We now prove that our parallel algorithm has low depth, and is work-efficient with respect to the sequential HDT algorithm. For simplicity, we assume that we start with no edges in a graph on n vertices.

THEOREM 3.4. *A batch of k edge deletions can be processed in $O(\lg^4 n)$ depth w.h.p.*

PROOF. The algorithm doubles the number of edges searched in each phase. Therefore, after $\lg m = O(\lg n)$ phases, all non-tree edges incident on the component will be searched.

In every round, each active component is either deactivated, or has a replacement edge found. In the worst case, the edges found for each active component pair the components off, leaving us with half as many active components in the subsequent round. As we lose a constant fraction of the active components per round, the algorithm takes $O(\lg n)$ rounds.

A given level can therefore perform at most $O(\lg^2 n)$ phases. Each phase consists of fetching, examining, and pushing down non-tree edges, and hence can be implemented in $O(\lg n)$ depth w.h.p. by Lemma 3.5, Theorem 2.4, and Lemma 3.6. Therefore, the overall depth for a given level is $O(\lg^3 n)$ w.h.p. As all $\lg n$ levels will be processed in the worst case, the overall depth of the algorithm is $O(\lg^4 n)$ w.h.p. \square

We now analyze the work performed by the algorithm. We begin by stating the following lemmas on the efficiency of the augmented ET-tree operations.

LEMMA 3.5. *Given some vertex, v in a batch-parallel ET-tree, we can fetch the first l tree (or non-tree) edges referenced by the augmented values in the tree in $O\left(l \lg\left(1 + \frac{n_c}{l}\right)\right)$ work and $O(\lg n)$ depth w.h.p. where n_c is the number of vertices in the ET-tree at the current level. Furthermore, removing the edges can be done in the same bounds.*

LEMMA 3.6. *Decreasing the level of l tree (or non-tree) edges in a batch-parallel ET-tree can be performed in $O\left(l \lg\left(1 + \frac{n_c}{l}\right)\right)$ expected work and $O(\lg n)$ depth w.h.p. where n_c is the number of nodes in the ET-tree at the current level.*

The proofs of these lemmas are provided in the full version of our paper [2]. We now have the tool to analyze the work performed by batch deletion.

LEMMA 3.7. *The work performed by BATCHDELETION excluding the calls to PARALLELEVESEARCH is*

$$O\left(k \lg n \lg\left(1 + \frac{n}{k}\right)\right),$$

in expectation.

PROOF. The edge deletions performed by Line 2 cost $O(k)$ work by Lemma 3.1. Filtering the tree edges (Line 3) can be done in $O(k)$ work. Deleting the tree edges costs at most $O(k \lg(1 + n/k))$ work by Lemma 2.3 (Line 4).

Line 5 perform a FINDREPR call for each endpoint of each deleted tree edge. These calls can be implemented as a single BATCHFINDREPR call which costs $O(k \lg(1 + n/k))$ work in expectation by Theorem 2.4. Since in the worst case each tree edge must be deleted from $\lg n$ levels, the overall cost of this step is $O(k \lg n \lg(1 + n/k))$

in expectation. Summing up the costs for each level proves the lemma. \square

THEOREM 3.8. *The expected amortized cost per edge insertion or deletion is $O(\lg^2 n)$.*

PROOF. Algorithm 3 takes as input a batch of k edge deletions. By Lemma 3.7, the expected work performed by `BATCHDELETION` excluding the calls to `PARALLELEVELSEARCH` is

$$O\left(k \lg n \lg\left(1 + \frac{n}{k}\right)\right),$$

which is at most $O(k \lg^2 n)$ in expectation. We now consider the cost of the calls to `PARALLELEVELSEARCH`. Specifically, we show that the work performed during the calls to `PARALLELEVELSEARCH` can either be charged to level decreases on edges, or is at most $O(k \lg n)$ per call in expectation. Since the total number of calls to `PARALLELEVELSEARCH` is at most $\lg n$, the bounds follow.

First, observe that the number of spanning forest edges we discover, $|S|$, is at most k , since at most k tree edges were deleted initially. Therefore, the batch insertion on Line 12 costs $O(k \lg n)$ in expectation by Theorem 2.4. Similarly, L , the number of components that are supplied to `PARALLELEVELSEARCH`, is at most k . Therefore, the cost of filtering the components in L based on their size, and checking whether their representative exists in F_i is at most $O(k \lg n)$ in expectation (Lines 13–14).

To fetch, examine, and push down l tree or non-tree edges costs

$$O\left(l \lg\left(1 + \frac{n}{l}\right)\right),$$

work in expectation, by Lemma 3.5, Theorem 2.4, and Lemma 3.6. Note that this is at most $O(\lg n)$ per edge. In particular, the cost of retrieving and pushing the tree edges of active components to level $i - 1$ (Line 6) is therefore at most $O(\lg n)$ per edge in expectation, which we charge to the corresponding level decreases.

We now show that all work done while searching for replacement edges (Lines 15–27) can be charged to level decreases. Consider an active component, c in some round. Suppose the algorithm performs $q > 0$ phases before either the component is exhausted (all incident non-tree edges have been checked), or a replacement edge is found. First consider the case where it finds a replacement edge. If $q = 1$, only a single edge was inspected, so then we charge the $\lg n$ work for the round to the edge, which will become a tree edge. Otherwise, it performs $q - 1$ phases which do not produce any replacement edge.

Since phase w inspects 2^w edges, it costs $O(2^w \lg n)$ work. The total work over all q phases is therefore

$$\sum_{w=0}^{q-1} 2^w \lg n = O(2^q \lg n)$$

in expectation. However, since no replacement was found during the first $q - 1$ phases, there are at least $2^{q-1} = O(2^q)$ edges that will be pushed down, so we can charge $O(\lg n)$ work to each such edge to pay for this. In the other case, q phases run without finding a replacement edge. In this case, all edges inspected are pushed down, and hence each assumes a cost of $O(\lg n)$ in expectation.

Now, we argue that the work done while processing the replacement edges is $O(k \lg n)$ in expectation over all rounds. Since k edges were deleted, the algorithm discovers at most k replacement edges.

We charge the work in these steps to the replacement edges that we find. Let k' be the number of replacement edges that we find. Filtering the edges, and computing a spanning forest all costs $O(k')$ work. Promoting the edges to tree edges (inserting them into F_i and updating the adjacency lists) costs $O(k' \lg n)$ work in expectation. Finally, updating the components costs $O(k' \lg n)$ work in expectation, which we can charge to either the component, if it is removed from C in this round, or to the replacement edge that it finds, which is promoted to a tree edge. Since the algorithm can find at most k replacement edges, the cost per level is $O(k \lg n)$ in expectation for these steps as necessary.

In total, on each level the algorithm performs $O(k \lg n)$ expected work that is not charged to a level decrease. Summing over $\lg n$ levels, this yields an amortized cost of $O(\lg^2 n)$ expected work per edge deletion. Finally, since the level of an edge can decrease at most $\lg n$ times, and an edge is charged $O(\lg n)$ expected work each time its level is decreased, the expected amortized cost per edge insertion is $O(\lg^2 n)$. \square

4 AN IMPROVED ALGORITHM

In this section we design an improved version of the parallel algorithm that performs less work than the algorithm from Section 3. Furthermore, the improved algorithm runs in $O(\lg^3 n)$ depth w.h.p., improving on Algorithm 4 which runs in $O(\lg^4 n)$ depth w.h.p.

4.1 The Interleaved Deletion Algorithm

Overview. Algorithm 5 is based on *interleaving* the phases of doubling that search for replacement edges with the spanning forest computation performed on the replacement edges. Recall that in Algorithm 4, the number of edges examined in each round *is reset*, and the doubling algorithm must therefore start with an initial search size of 1 on the next round. Because the doubling resets from round to round, the number of phases per round can be $O(\lg n)$ in the worst case, making the total number of phases per level $O(\lg^2 n)$. Instead, the interleaved algorithm avoids resetting the search size by maintaining a *single*, geometrically increasing search size over all rounds of the search.

The second important difference in Algorithm 5 compared with Algorithm 4 is that it *defers* inserting tree edges found on this level until the end of the search. Instead, it continues to search for replacement edges from the *initial* components until the component is deactivated. This property is important to show that the work done for a component across all rounds is dominated by the cost of the last round, since the number of vertices in the component is fixed, but the number of non-tree edges examined doubles in each round. For the same reason, it also defers inserting the pushed edges onto level $i - 1$. We crucially use this property to obtain improved batch work bounds (Section 4.3).

Another difference in the modified algorithm is that if a component is still active after adding the replacement edges found in this round (i.e., the component on level i still has size at most 2^{i-1}), then *all* of the edges found in this round can be pushed to level $i - 1$ without violating Invariant 1. Notice now that when pushing down edges, both the *tree and non-tree* edges that are found in this round

Algorithm 5 The interleaved level search algorithm

```
1: procedure COMPONENTSEARCH( $i, c, s$ )
2:    $w_{\max} \leftarrow c.\text{NUMNONTREEEDGES}$ 
3:    $w \leftarrow \min(s, w_{\max})$ 
4:    $E_c \leftarrow$  First  $w$  non-tree edges in  $c$ 
5:   return {All replacement edges in  $E_c$ }
6: procedure PUSHEDGES( $i, c, s, M$ )
7:    $w_{\max} \leftarrow c.\text{NUMNONTREEEDGES}$ 
8:    $w \leftarrow \min(s, w_{\max})$ 
9:    $E_c \leftarrow$  {First  $w$  non-tree edges in  $c$ }
10:  if  $M[c].\text{SIZE} \leq 2^{i-1}$  and  $w < w_{\max}$  then
11:    Remove edges in  $E_c$  from level  $i$ 
12:    return  $E_c$ 
13:  return  $\emptyset$ 
14: procedure INTERLEAVEDLEVELSEARCH( $i, L = \{c_1, c_2, \dots\}, S$ )
15:   $F_i.\text{BATCHINSERT}(S)$ 
16:   $C \leftarrow c \in L$  with  $\text{size} \leq 2^{i-1}$ 
17:   $D \leftarrow c \in L$  with  $\text{size} > 2^{i-1}$ 
18:  Push level  $i$  tree edges of all components in  $C$  to level  $i - 1$ 
19:   $r \leftarrow 0, T \leftarrow \emptyset, E_P \leftarrow \emptyset$ 
20:   $M \leftarrow \{c \rightarrow c \mid c \in C\}$ 
21:  while  $|C| > 0$  do
22:     $w \leftarrow 2^r$ 
23:     $R \leftarrow \cup_{c \in C} \text{COMPONENTSEARCH}(i, c, w)$   $\triangleright$  In parallel
24:     $R' \leftarrow \{(F_i.\text{FINDREPR}(u), F_i.\text{FINDREPR}(v)) \mid (u, v) \in R\}$ 
25:     $T'_r \leftarrow \text{SPANNINGFOREST}(R')$ 
26:     $T_r \leftarrow$  Edges in  $R$  corresponding to edges in  $T'_r$ 
27:     $T \leftarrow T \cup T_r$ 
28:    Update  $M$ , the map of supercomponents and their sizes
29:     $E_P \leftarrow E_P \cup_{c \in C} \text{PUSHEDGES}(i, c, w, M)$   $\triangleright$  In parallel
30:     $D_r \leftarrow \{c \in C$  with no non-tree edges, or  $\text{size} > 2^{i-1}\}$ 
31:     $D \leftarrow D \cup D_r$ 
32:     $C \leftarrow C \setminus D_r$ 
33:     $r \leftarrow r + 1$ 
34:  Promote edges in  $T$  not in  $E_P$  to tree edges at level  $i$ 
35:   $F_i.\text{BATCHINSERT}(T)$ 
36:  Insert non-tree and tree edges in  $E_P$  to level  $i - 1$ 
37:  return  $(D, S \cup T)$ 
```

are pushed. Pushing down all edges ensures that the algorithm performs enough level decreases to which to charge the work performed during the next round. The component deactivates either once it runs out of incident non-tree edges, or when it becomes too large. Since the algorithm defers adding the new tree edges found until the end of the level, it also maintains an auxiliary data structure that dynamically tracks the size of the resulting components as new edges are found.

The Deletion Algorithm. We briefly describe the main differences between INTERLEAVEDLEVELSEARCH, the new level search procedure, and PARALLELEVELSEARCH. The algorithm consists of a number of **rounds** (Lines 21–33). We use r to track the round numbers, and we use E_P to store the set of both tree and non-tree edges that will be pushed to level $i - 1$ at the end of the search at this level (Line 19). T stores the set of tree edges that have been selected, which will be added to the spanning forest at the end of the level. Lastly, we use M to maintain a dynamic mapping from all

the components in L to a unique representative for their contracted supercomponent (initially itself), and the size of the contracted supercomponent.

In round r , the algorithm first retrieves the first 2^r (or fewer) edges of each the active components in parallel, and finds replacement edges. All replacement edges are added to the set R (line 23).

The algorithm then computes a spanning forest over the edges in R , and computes T_r , which are the original replacement edges in R that were selected as spanning forest edges (lines 25–27). The spanning forest computation returns, in addition to the tree edges, a mapping from the vertices in R' to their connectivity label (line 25), which can be used on line 28 to efficiently update the representatives of all affected components and the sizes of the supercomponents.

The next step maps over the components in parallel again, calling PUSHEDGES on each active component, and checks whether the edges searched in this round can be (lazily) pushed to level $i - 1$ (Line 29).³ If a component is still active (its new size is small enough to still be searched, and the component still has some non-tree edges remaining) (line 10), all of the searched edges are removed from the adjacency lists at level i (line 11) and are added to the set of edges that will be pushed to level $i - 1$ at the end of the level (Lines 12 and 29). Note that this set of edges contains both replacement tree edges we discovered, and non-tree edges. The tree-edges can be pushed down to level $i - 1$ because the component with the tree edges added has size $\leq 2^{i-1}$.

The end of the round (lines 30–33) handles updating the set of components and incrementing the round number, as in Algorithm 4.

Finally, once all components are inactive, the tree edges found at this level that are not contained in E_P are promoted (the tree edges added to E_P have their level decreased to $i - 1$) and inserted into F_i (Lines 34–35), and all edges added to E_P in Line 29 are pushed down to level $i - 1$ (Line 36). Note that any tree-edges found in this set are promoted in level $i - 1$ and added to F_{i-1} . The procedure returns the set of components and all replacement edges found at this level and levels below it (Line 37).

4.2 Cost Bounds

We start by showing that the depth of Algorithm 5 is $O(\lg^3 n)$.

LEMMA 4.1. *The number of rounds performed by Algorithm 5 is $O(\lg n)$ and the depth of each round is $O(\lg n)$ w.h.p.. The depth of the INTERLEAVEDLEVELSEARCH is therefore $O(\lg^2 n)$ w.h.p..*

PROOF. Each round of the algorithm increases the search size of a component by a factor of 2. Therefore, after $O(\lg n)$ rounds, every non-tree edge incident on a component will be considered and the algorithm will terminate.

To argue the depth bound, we consider the main steps performed during a round. Fetching, examining and removing the edges from level i takes $O(\lg n)$ depth w.h.p. by Lemma 3.5, Theorem 2.4, and Lemma 3.6. Computing a spanning forest on the replacement edges and filtering the components (at most k replacement edges, or components) can be done in $O(\lg k)$ depth. The depth per round is

³ Note that the set of edges retrieved by PUSHEDGES in Line 9 is assumed to be the same as the one in Line 4. This assumption is satisfied by using our FETCHEDGES primitive on a batch-parallel ET-tree, and can be satisfied in general by associating the edges retrieved in COMPONENTSEARCH to be used in PUSHEDGES.

therefore $O(\lg n)$ w.h.p. and the depth of `INTERLEAVEDLEVELSEARCH` is $O(\lg^2 n)$ w.h.p. \square

Combining Lemma 4.1 with the fact that there are $\lg n$ levels gives the following theorem.

THEOREM 4.2. *A batch of k edge deletions can be processed in $O(\lg^3 n)$ depth w.h.p.*

We now consider the work performed by the algorithm. We start with a lemma showing that the search-size for a component increases geometrically until the round where the component is deactivated.

LEMMA 4.3. *Consider a component, c , that is active at the end of round $r - 1$. If c is not removed from C , then it examines $\geq 2^{r-1}$ edges that are pushed down to level $i - 1$ at the end of the search.*

PROOF. We prove the contrapositive. Suppose that $< 2^{r-1}$ edges are pushed down in total by c in the last round. Then, we will show that c cannot be active in the next round (i.e., it is removed from C in round $r - 1$).

Notice that c must be active at the start of round $r - 1$. Consider the check on Line 10, which checks whether $w \leq 2^{r-1}$ and $w < w_{\max}$ on this round. Suppose for the sake of contradiction that both conditions are true. Then, by the fact that $w < w_{\max}$, it must be the case that $w = 2^{r-1}$ by Line 8. If the condition is true, then on Line 11 the algorithm adds 2^{r-1} edges to be pushed to level $i - 1$, contradicting our assumption that $< 2^{r-1}$ edges are pushed.

Therefore the check on Line 10 must be false, giving that either $w > 2^{i-1}$, or $w = w_{\max}$. This means that c will be marked as inactive on Line 30, and then become deactivated on line 32. Therefore, if $< 2^{r-1}$ edges are pushed down by c in round $r - 1$, c is deactivated at the end of the round, concluding the proof. \square

LEMMA 4.4. *Consider the work done by some component c over the course of `INTERLEAVEDLEVELSEARCH` at a given level. Let R be the total number rounds that c is active. Then, c pushes down $p_c = 2^R - 1$ edges in total. Furthermore, the total cost of searching for and pushing down replacement edges performed by c is*

$$O\left(p_c \lg\left(1 + \frac{n_c}{p_c}\right)\right)$$

in expectation, where n_c is the number of vertices in c .

PROOF. By Lemma 4.3, for each round $r < R$, c adds 2^r edges to be pushed down. Summing over all rounds shows that the total number of edges added to be pushed down is $2^R - 1$. The cost of pushing down these edges at the end of the search at this level is exactly

$$O\left(p_c \lg\left(1 + \frac{n_c}{p_c}\right)\right).$$

by Lemma 3.6, since the size of the tree that is affected is n_c .

We now consider the cost of fetching and examining the edges over all rounds. The cost of fetching and examining 2^r edges is

$$O\left(2^r \lg\left(1 + \frac{n_c}{2^r}\right)\right),$$

in expectation by Theorem 2.4 and Lemma 3.5. Summing over all rounds $r < R$, the work is

$$\sum_{r=1}^{R-1} O\left(2^r \lg\left(1 + \frac{n_c}{2^r}\right)\right)$$

in expectation to fetch and examine edges in the first $R - 1$ rounds, which is equal to

$$O\left(2^R \lg\left(1 + \frac{n_c}{2^R}\right)\right),$$

by Lemma 2.2. Since on round R , the algorithm searches at most 2^R edges, the total cost of searching for replacement edges over all rounds is at most

$$O\left(2^R \lg\left(1 + \frac{n_c}{2^R}\right)\right) = O\left(p_c \lg\left(1 + \frac{n_c}{p_c}\right)\right).$$

\square

LEMMA 4.5. *The cost of `INTERLEAVEDLEVELSEARCH` is at most*

$$O\left(k \lg\left(1 + \frac{n}{k}\right) + p \lg\left(1 + \frac{n}{p}\right)\right)$$

in expectation where p is the total number of edges pushed down.

PROOF. First consider lines 2–5. Since we are deleting a batch of k edges, we can find at most k replacement edges to reconnect these components. Therefore line 2 performs $O\left(k \lg\left(1 + \frac{n}{k}\right)\right)$ expected work by Theorem 2.4. Pushing t spanning tree edges to the next level (line 5) can be done in $O\left(t \lg\left(\frac{n}{t} + 1\right)\right)$ expected work by Lemmas 3.5, 3.6, and 2.1, and Theorem 2.4. Hence in total, lines 2–5 perform at most $O\left(k \lg\left(1 + \frac{n}{k}\right) + t \lg\left(1 + \frac{n}{t}\right)\right)$ work in expectation.

Now, consider the cost of the steps which scan or update the components that are active in each round. On the first round, this cost is $O(k)$. In every subsequent round, r , by Lemma 4.3 each currently active component must have added 2^{r-1} edges to be pushed down on the previous round. Therefore, we can charge the $O(1)$ work per component performed in this round to these edge pushes.

Next, we analyze the work done while searching for and pushing replacement edges. Consider some component $c \in C$ that is searched on this level. By Lemma 4.4, the cost of searching for and pushing down the replacement edges incident on this component is

$$O\left(p_c \lg\left(1 + \frac{n_c}{p_c}\right)\right)$$

in expectation, where n_c is the number of vertices in c and p_c is the total number of edges pushed down by c .

The total work done over all components to search for replacement edges and push down both the original tree edges, and the edges in each round is therefore

$$O\left(t \lg\left(1 + \frac{n}{t}\right) + \sum_{c \in C} p_c \lg\left(1 + \frac{n_c}{p_c}\right)\right).$$

in expectation. Since $\sum n_c = n$, by Lemma 2.1 this costs

$$O\left(p \lg\left(1 + \frac{2n}{p}\right)\right) = O\left(p \lg\left(1 + \frac{n}{p}\right)\right)$$

work in expectation, where $p = t + \sum p_c$ is the total number of edges pushed, including tree and non-tree edges. Therefore, the total cost is

$$O\left(k \lg\left(1 + \frac{n}{k}\right) + p \lg\left(1 + \frac{n}{p}\right)\right)$$

in expectation. \square

THEOREM 4.6. *The expected amortized cost per edge insertion or deletion is $O(\lg^2 n)$.*

PROOF. The proof follows from the same argument as Theorem 3.8, by using Lemma 4.5. \square

4.3 A Better Work Bound

We now show that by a more careful analysis, we can obtain a tighter bound on the amount of work performed by the interleaved algorithm. In particular, we show in this section that the algorithm performs

$$O\left(\lg n \lg\left(1 + \frac{n}{\Delta}\right)\right)$$

amortized work per edge in expectation, where Δ is the average batch size of all batches of deletions. Therefore, if we process batches of deletions of size $O(n/\text{polylog}(n))$ on average, our algorithm performs $O(\lg n \lg \lg n)$ expected amortized work per edge, rather than $O(\lg^2 n)$. Furthermore, if we have batches of size $O(n)$, the cost is just $O(\lg n)$ per edge.

At a high level, our proof formalizes the intuition that in the worst case, all edges are pushed down at every level, and that performing fewer deletion operations results in larger batches of pushes which take advantage of work bounds of the ET-tree. Our proof crucially relies on the fact that although the deletion algorithm at a level can perform $O(\lg n)$ ET-tree operations per component, since the batch sizes are geometrically increasing, these operations have the cost of a single ET-tree operation per component. Furthermore, Lemma 4.5 shows that the costs per component can be combined so that the total cost is equivalent to the cost of a single ET-tree operation on all the vertices. Therefore, the number of deletion operations can be exactly related to the effective number of ET-tree operations at a level. We relate the number of deletions to the average batch size, which lets us obtain a single unified bound for both insertions and deletions.

THEOREM 4.7. *Using the interleaved deletion algorithm, the amortized work performed by BATCHDELETION and BATCHINSERTION on a batch of k edges is*

$$O\left(k \lg n \lg\left(1 + \frac{n}{\Delta}\right)\right),$$

in expectation where Δ is the average batch size of all batch deletions.

PROOF. Batch insertions perform only $O\left(k \lg\left(1 + \frac{n}{k}\right)\right)$ work by Theorem 3.3, so we focus on the cost of deletion since it dominates. Consider the total amount of work performed by all batch deletion operations at any given point in the lifetime of the data structure. We will denote by k_b , the size of batch b , and by $p_{b,i}$, the number of edges pushed down on level i during batch b . Combining Lemmas 3.7, and 4.5, the total work is bounded above by

$$O\left(\sum_{\text{batch } b} \sum_{\text{level } i} k_b \lg\left(1 + \frac{n}{k_b}\right) + p_{b,i} \lg\left(1 + \frac{n}{p_{b,i}}\right)\right).$$

We begin by analyzing the first term, which is paid for by the deletion algorithm. Let

$$K = \sum_{\text{batch } b} k_b$$

denote the total number of deleted edges. Applying Lemma 2.1, and using the fact that there are $\lg n$ levels, we have

$$O\left(\sum_{\text{batch } b} \sum_{\text{level } i} k_b \lg\left(1 + \frac{n}{k_b}\right)\right) = O\left(K \lg n \lg\left(1 + \frac{n \cdot d}{K}\right)\right),$$

where d is the number of batches of deletions. Since $K/d = \Delta$, this is equal to

$$O\left(K \lg n \lg\left(1 + \frac{n}{\Delta}\right)\right),$$

work in expectation. Each batch can therefore be charged a cost of $\lg n \lg(1 + n/\Delta)$ per edge, and hence the amortized cost of batch deletion is

$$O\left(k \lg n \lg\left(1 + \frac{n}{\Delta}\right)\right)$$

in expectation.

The remainder of the cost, which comes entirely from searching for replacement edges, is charged to the insertions. Consider this cost and let

$$P = \sum_{\text{batch } b} \sum_{\text{level } i} p_{b,i}$$

denote the total such number of edge pushes. Since the total number of terms in the double sum is $d \lg n$, Lemma 2.1 allows us to bound the total work of all pushes by

$$\sum_{\text{batch } b} \sum_{\text{level } i} p_{b,i} \lg\left(1 + \frac{n}{p_{b,i}}\right) = O\left(P \lg\left(1 + \frac{nd \lg n}{P}\right)\right).$$

in expectation. Since every edge can only be pushed down once per level, we have

$$P \leq m \lg n,$$

where m is the total number of edges ever inserted. Therefore by Lemma 2.3, the total work is at most

$$O\left(m \lg n \lg\left(1 + \frac{nd \lg n}{m \lg n}\right)\right) = O\left(m \lg n \lg\left(1 + \frac{nd}{m}\right)\right)$$

in expectation. Since $d = K/\Delta$, this is equal to

$$O\left(m \lg n \lg\left(1 + \frac{nK}{m\Delta}\right)\right)$$

in expectation. Since each edge can be deleted only once, we have $K \leq m$, and hence we obtain that the total work to push all tree edges down is at most

$$O\left(m \lg n \lg\left(1 + \frac{n}{\Delta}\right)\right).$$

in expectation. We can therefore charge $O(\lg n \lg(1 + n/\Delta))$ per edge to each batch insertion. Since this dominates the cost of the insertion algorithm itself, the amortized cost of batch insertion is therefore

$$O\left(k \lg n \lg\left(1 + \frac{n}{\Delta}\right)\right),$$

in expectation as desired, concluding the proof. \square

5 RELATED WORK

Parallel Dynamic Algorithms. There are only a few results on parallel dynamic algorithms. Earlier results [20, 23] are not work-efficient with respect to the fastest sequential dynamic algorithms, do not support batch updates, and perform polynomial work per update. Some more recent results such as parallel dynamic depth-first search [41] and minimum spanning forest [42] process updates one at a time, and are therefore not batch-dynamic algorithms. Work efficient parallel batch-dynamic algorithms include those for the well-spaced point sets problem [3] and those for the dynamic trees problem [1, 52, 65].

Parallel Connectivity. Parallel algorithms for connectivity have a long history [10, 19, 32, 39, 48, 50, 56, 66], and there are many existing algorithms that solve the problem work-efficiently and in low-depth [18, 24, 27, 28, 47, 49, 58], some of which are also practical [21, 58]. However, there is no obvious way to adapt existing parallel connectivity algorithms to the dynamic setting, particularly for batch updates.

Parallel Dictionaries and Trees. There are many results on parallel dictionaries and trees supporting batch updates [7, 13, 15, 16, 25, 57, 60]. The dictionary data structures in the literature culminated in dictionaries supporting batch insertions, deletions and lookups in linear work and $O(\lg^* n)$ depth w.h.p. [25]. Early work on batch insertions into trees focused on optimizing the depth, but was not work-efficient. Paul et al. design batch search, insertion and deletion algorithms for 2-3 trees on the EREW PRAM [46]. These results were later extended to B-trees by Higham et al. [31]. The algorithms of both Paul et al. and Higham et al. perform $O(m \lg n)$ work for m tree operations.

Recent work on parallel tree data structures has focused on how to parallelize batch operations for various balancing schemes in binary search trees [13], and also how to improve the depth of these operations [7]. There is also some very recent work on extending these tree data structures to support range and segment queries [60] as well as practical implementations of parallel trees supporting batch insertions, deletions and lookups [61].

Other Related Work. There is also recent work on parallel working-set structures that supports batching by Agrawal et al. [5]. Earlier work by Agrawal et al. [4] introduces the idea of implicit batching which uses scheduler support to convert dynamically multithreaded programs using an abstract data type to programs that perform batch accesses to an underlying parallel data structure.

6 DISCUSSION

In this paper, we present a novel batch-dynamic algorithm for the connectivity problem. Our algorithm is always work-efficient with respect to the Holm, de Lichtenberg and Thorup dynamic connectivity algorithm, and is asymptotically faster than their algorithm when the average batch size is sufficiently large. A parallel implementation of our algorithm achieves $O(\lg^3 n)$ depth w.h.p., and is, to the best of our knowledge, the first parallel algorithm for the dynamic connectivity problem performing $O(T \text{ polylog}(n))$ total expected work, where T is the total number of edge operations.

There are several natural questions to address in future work. First, can the depth of our algorithm be improved to $O(\lg^2 n)$ without increasing the work? Investigating lower bounds in the batch setting would also be very interesting—are there non-trivial lower-bounds for batch-dynamic connectivity? Lastly, in this paper we show expected amortized bounds. One approach to strengthen these bounds is to show that our tree operations hold w.h.p. and argue that our amortized bounds hold w.h.p. Another is to design a deterministic batch-dynamic forest connectivity data structure with the same asymptotic complexity as the batch-parallel ET-tree, which would make the randomized bounds in this paper deterministic.

Two additional questions are whether we can extend our results to give parallel work-efficient batch-dynamic MST, 2-edge connectivity and biconnectivity algorithms. MST seems solvable using the techniques presented in this paper, although our dynamic tree structure would need to be extended with additional primitives. Existing sequential 2-edge connectivity and biconnectivity algorithms require a dynamic tree data structure supporting path queries which are not supported by ET-trees. However, RC-trees [1] can be extended to support path queries, which makes them a possible candidate for this line of work. Finally, it seems likely that ideas from our work can be extended to give a parallel batch-dynamic Monte-Carlo connectivity algorithm based on the Kapron-King-Mountjoy algorithm [38].

ACKNOWLEDGEMENTS

We thank Tom Tseng and Goran Zuzic for helpful discussions. This work was supported in part by NSF grants CCF-1408940, CCF-1533858, and CCF-1629444.

REFERENCES

- [1] U. A. Acar, V. Aksenov, and S. Westrick. Brief announcement: Parallel dynamic tree contraction via self-adjusting computation. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2017.
- [2] U. A. Acar, D. Anderson, G. E. Blelloch, and L. Dhulipala. Parallel batch-dynamic graph connectivity. *arXiv preprint arXiv:1903.08794*, 2019.
- [3] U. A. Acar, A. Cotter, B. Hudson, and D. Türkoglu. Parallelism in dynamic well-spaced point sets. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2011.
- [4] K. Agrawal, J. T. Fineman, K. Lu, B. Sheridan, J. Sukha, and R. Utterback. Provably good scheduling for parallel programs that use data structures through implicit batching. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2014.
- [5] K. Agrawal, S. Gilbert, and W. Q. Lim. Parallel working-set search structures. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2018.
- [6] K. J. Ahn, S. Guha, and A. McGregor. Analyzing graph structure via linear measurements. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 459–467, 2012.
- [7] Y. Aikhrentsev and P. Sanders. Fast parallel operations on search trees. In *IEEE International Conference on High Performance Computing (HiPC)*, 2016.
- [8] A. Andoni, C. Stein, Z. Song, Z. Wang, and P. Zhong. Parallel graph connectivity in log diameter rounds. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2018.
- [9] N. S. Arora, R. D. Blumofe, and C. G. Plaxton. Thread scheduling for multi-programmed multiprocessors. *Theory of Computing Systems (TOCS)*, 34(2), Apr 2001.
- [10] B. Awerbuch and Y. Shiloach. New connectivity and MSF algorithms for Ultra-computer and PRAM. In *International Conference on Parallel Processing (ICPP)*, 1983.
- [11] G. E. Blelloch. Programming parallel algorithms. *Commun. ACM*, 39(3), Mar. 1996.
- [12] G. E. Blelloch and L. Dhulipala. Introduction to parallel algorithms. <http://www.cs.cmu.edu/~realworld/slidesS18/parallelChap.pdf>, 2018. Carnegie Mellon University.
- [13] G. E. Blelloch, D. Ferizovic, and Y. Sun. Just join for parallel ordered sets. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2016.

- [14] G. E. Blelloch, P. B. Gibbons, and Y. Matias. Provably efficient scheduling for languages with fine-grained parallelism. *J. ACM*, 46(2), Mar. 1999.
- [15] G. E. Blelloch and M. Reid-Miller. Fast set operations using treaps. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 1998.
- [16] G. E. Blelloch and M. Reid-Miller. Pipelining with futures. *Theory of Computing Systems (TOCS)*, 32(3), 1999.
- [17] R. D. Blumofe and C. E. Leiserson. Scheduling multithreaded computations by work stealing. *J. ACM*, 46(5):720–748, 1999.
- [18] R. Cole, P. N. Klein, and R. E. Tarjan. Finding minimum spanning forests in logarithmic time and linear work using random sampling. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 1996.
- [19] R. Cole and U. Vishkin. Approximate parallel scheduling. II. Applications to logarithmic-time optimal parallel graph algorithms. *Information and Computation*, 92(1):1–47, 1991.
- [20] S. K. Das and P. Ferragina. An $o(n)$ work EREW parallel algorithm for updating MST. In *European Symposium on Algorithms (ESA)*, 1994.
- [21] L. Dhulipala, G. E. Blelloch, and J. Shun. Theoretically efficient parallel graph algorithms can be fast and scalable. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2018.
- [22] D. Eppstein, Z. Galil, G. F. Italiano, and A. Nissenzweig. Sparsification—a technique for speeding up dynamic graph algorithms. *J. ACM*, 44(5):669–696, 1997.
- [23] P. Ferragina and F. Luccio. Batch dynamic algorithms for two graph problems. *International Conference on Parallel Architectures and Languages Europe (PARLE)*, 1994.
- [24] H. Gazit. An optimal randomized parallel algorithm for finding connected components in a graph. *SIAM J. on Computing*, 20(6):1046–1067, 1991.
- [25] J. Gil, Y. Matias, and U. Vishkin. Towards a theory of nearly constant time parallel algorithms. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, 1991.
- [26] Y. Gu, J. Shun, Y. Sun, and G. E. Blelloch. A top-down parallel semisor. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2015.
- [27] S. Halperin and U. Zwick. An optimal randomized logarithmic time connectivity algorithm for the EREW PRAM (extended abstract). In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 1994.
- [28] S. Halperin and U. Zwick. Optimal randomized EREW PRAM algorithms for finding spanning forests. In *J. Algorithms*, 2000.
- [29] M. R. Henzinger and V. King. Randomized dynamic graph algorithms with polylogarithmic time per operation. In *ACM Symposium on Theory of Computing (STOC)*, ACM, 1995.
- [30] M. R. Henzinger and V. King. Maintaining minimum spanning forests in dynamic graphs. *SIAM J. on Computing*, 31(2):364–374, 2001.
- [31] L. Higham and E. Schenck. Maintaining B-trees on an EREW PRAM. *J. Parallel Distrib. Comput.*, 22(2), 1994.
- [32] D. S. Hirschberg, A. K. Chandra, and D. V. Sarwate. Computing connected components on parallel computers. *Commun. ACM*, 22(8):461–464, Aug. 1979.
- [33] J. Holm, K. De Lichtenberg, and M. Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48(4):723–760, 2001.
- [34] S.-E. Huang, D. Huang, T. Kopelowitz, and S. Pettie. Fully dynamic connectivity in $O(\log n(\log \log n)^2)$ amortized expected time. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 510–520, 2017.
- [35] A. Iyer, L. E. Li, and I. Stoica. CellIQ: Real-time cellular network analytics at scale. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2015.
- [36] A. P. Iyer, L. E. Li, T. Das, and I. Stoica. Time-evolving graph processing at scale. In *International Workshop on Graph Data Management Experiences and Systems (GRADES)*, 2016.
- [37] J. JaJa. *Introduction to Parallel Algorithms*. Addison-Wesley Professional, 1992.
- [38] B. M. Kapron, V. King, and B. Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013.
- [39] D. R. Karger, N. Nisan, and M. Parnas. Fast connected components algorithms for the EREW PRAM. *SIAM J. on Computing*, 28(3):1021–1034, Feb. 1999.
- [40] C. Kejlberg-Rasmussen, T. Kopelowitz, S. Pettie, and M. Thorup. Faster Worst Case Deterministic Dynamic Connectivity. In *European Symposium on Algorithms (ESA)*, 2016.
- [41] S. Khan. Near optimal parallel algorithms for dynamic DFS in undirected graphs. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2017.
- [42] T. Kopelowitz, E. Porat, and Y. Rosenmutter. Improved worst-case deterministic parallel dynamic minimum spanning forest. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2018.
- [43] R. McColl, O. Green, and D. A. Bader. A new parallel algorithm for connected components in dynamic graphs. In *IEEE International Conference on High Performance Computing (HiPC)*, 2013.
- [44] P. B. Miltersen, S. Subramanian, J. S. Vitter, and R. Tamassia. Complexity models for incremental computation. *Theoretical Computer Science (TCS)*, 130(1), 1994.
- [45] D. Nanongkai and T. Saranurak. Dynamic spanning forest with worst-case update time: adaptive, las vegas, and $O(n^{1/2-\epsilon})$ -time. In *ACM Symposium on Theory of Computing (STOC)*, ACM, 2017.
- [46] W. Paul, U. Vishkin, and H. Wagener. Parallel dictionaries on 2–3 trees. In *Intl. Colloq. on Automata, Languages and Programming (ICALP)*, 1983.
- [47] S. Pettie and V. Ramachandran. A randomized time-work optimal parallel algorithm for finding a minimum spanning forest. *SIAM J. on Computing*, 31(6), 2002.
- [48] C. A. Phillips. Parallel graph contraction. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 1989.
- [49] C. K. Poon and V. Ramachandran. A randomized linear work EREW PRAM algorithm to find a minimum spanning forest. In *International Symposium on Algorithms and Computation (ISAAC)*, 1997.
- [50] J. Reif. Optimal parallel algorithms for integer sorting and graph connectivity. *TR-08-85, Harvard University*, 1985.
- [51] J. H. Reif and S. Sen. Parallel computational geometry: An approach using randomization. In J. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, chapter 18. Elsevier Science, 1999.
- [52] J. H. Reif and S. R. Tate. Dynamic parallel tree contraction (extended abstract). In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 1994.
- [53] O. Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4):17:1–17:24, Sept. 2008.
- [54] S. Sahu, A. Mhedhbi, S. Salioglu, J. Lin, and M. T. Özsu. The ubiquity of large graphs and surprising challenges of graph processing. *Proceedings of the VLDB Endowment (PVLDB)*, 11(4):420–431, 2017.
- [55] M. Sha, Y. Li, B. He, and K.-L. Tan. Accelerating dynamic graph analytics on GPUs. *Proceedings of the VLDB Endowment (PVLDB)*, 11(1):107–120, Sept. 2017.
- [56] Y. Shiloach and U. Vishkin. An $O(\lg n)$ parallel connectivity algorithm. *J. Algorithms*, 3(1):57–67, 1982.
- [57] J. Shun and G. E. Blelloch. Phase-concurrent hash tables for determinism. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2014.
- [58] J. Shun, L. Dhulipala, and G. Blelloch. A simple and practical linear-work parallel algorithm for connectivity. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2014.
- [59] N. Simiri, K. Tangwongsan, S. Tirthapura, and K.-L. Wu. Work-efficient parallel union-find with applications to incremental graph connectivity. In *European Conference on Parallel Processing (Euro-Par)*, 2016.
- [60] Y. Sun and G. E. Blelloch. Parallel range and segment queries with augmented maps. *arXiv preprint:1803.08621*, 2018.
- [61] Y. Sun, D. Ferizovic, and G. E. Blelloch. Pam: Parallel augmented maps. In *ACM Symposium on Principles and Practice of Parallel Programming (PPoPP)*, 2018.
- [62] R. E. Tarjan. Efficiency of a good but not linear set union algorithm. *J. ACM*, 22(2):215–225, 1975.
- [63] M. Thorup. Decremental dynamic connectivity. *J. Algorithms*, 33(2):229–243, 1999.
- [64] M. Thorup. Near-optimal fully-dynamic graph connectivity. In *ACM Symposium on Theory of Computing (STOC)*, ACM, 2000.
- [65] T. Tseng, L. Dhulipala, and G. Blelloch. Batch-parallel Euler tour trees. *Algorithm Engineering and Experiments (ALENEX)*, 2019.
- [66] U. Vishkin. An optimal parallel connectivity algorithm. *Discrete Applied Mathematics*, 9(2):197–207, 1984.
- [67] K. Vora, R. Gupta, and G. Xu. KickStarter: Fast and accurate computations on streaming graphs via trimmed approximations. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2017.
- [68] C. Wickramaarachchi, A. Kumbhare, M. Frincu, C. Chelms, and V. K. Prasanna. Real-time analytics for fast evolving social graphs. In *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, 2015.
- [69] C. Wulff-Nilsen. Faster deterministic fully-dynamic graph connectivity. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013.
- [70] C. Wulff-Nilsen. Fully-dynamic minimum spanning forest with improved worst-case update time. In *ACM Symposium on Theory of Computing (STOC)*, ACM, 2017.