A Deterministic Distributed Algorithm for Exact Weighted All-Pairs Shortest Paths in $\tilde{O}(n^{3/2})$ Rounds

Udit Agarwal*

Dept. of Computer Science,
University of Texas at Austin
Austin, TX, USA
udit@cs.utexas.edu

Valerie King[†]

Dept. of Computer Science,
University of Victoria
Victoria, BC, Canada
val@uvic.ca

Vijaya Ramachandran*

Dept. of Computer Science,
University of Texas at Austin
Austin, TX, USA
vlr@cs.utexas.edu

Matteo Pontecorvi

Nokia Bell Labs
Nozay, Île-de-France, France
matteo.pontecorvi@nokia.com

ABSTRACT

We present a deterministic distributed algorithm to compute all-pairs shortest paths (APSP) in an edge-weighted directed or undirected graph. Our algorithm runs in $\tilde{O}(n^{3/2})$ rounds in the Congest model, where n is the number of nodes in the graph. This is the first $o(n^2)$ rounds deterministic distributed algorithm for the weighted APSP problem. Our algorithm is fairly simple and incorporates a deterministic distributed algorithm we develop for computing a 'blocker set' [14], which has been used earlier in sequential dynamic computation of APSP.

CCS CONCEPTS

• **Theory of computation** → **Shortest paths; Distributed algorithms;**

KEYWORDS

Shortest paths; Distributed Algorithms; Blocker set; CONGEST Model

ACM Reference Format:

Udit Agarwal, Vijaya Ramachandran, Valerie King, and Matteo Pontecorvi. 2018. A Deterministic Distributed Algorithm for Exact Weighted All-Pairs Shortest Paths in $\tilde{O}(n^{3/2})$ Rounds. In *PODC '18: ACM Symposium on Principles of Distributed Computing, July 23-27, 2018, Egham, United Kingdom*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3212734.3212773>

*This work was supported by NSF Grant CCF-1320675.

[†]This work was supported by an NSERC Discovery Grant and the Simons Institute.

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.

PODC '18, July 23-27, 2018, Egham, United Kingdom

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5795-1/18/07...\$15.00

<https://doi.org/10.1145/3212734.3212773>

1 INTRODUCTION

The design of distributed algorithms for various network (or graph) problems such as shortest paths [3, 13, 17, 18] and minimum spanning tree [7, 8, 15, 21] is a well-studied area of research. The most widely considered model for studying distributed algorithms is the CONGEST model [19] (also see [3, 9, 12, 13, 17, 18]), described in more detail below. In this paper we consider the problem of computing all pairs shortest paths (APSP) in a weighted directed (or undirected) graph in this model.

The problem of computing all pairs shortest paths (APSP) in distributed networks is a very fundamental problem, and there has been a considerable line of work for the CONGEST model as described below in Section 1.2. However, for a weighted graph no deterministic algorithm was known in this model other than a trivial method that runs in n^2 rounds. In this paper we present the first algorithm for this problem in the CONGEST model that computes weighted APSP deterministically in less than n^2 rounds. Our algorithm computes APSP deterministically in $O(n^{3/2} \cdot \sqrt{\log n})$ rounds in this model in both directed and undirected graphs.

Our distributed APSP algorithm is quite simple and we give an overview in Section 2. It uses the notion of a blocker set introduced by King [14] in the context of sequential fully dynamic APSP computation. Our deterministic distributed algorithm for computing a blocker set is the most nontrivial component of our algorithm, and is described in Section 3.

1.1 The CONGEST Model

In the CONGEST model [19], there are n independent processors interconnected in a network. We refer to these processors as nodes. These nodes are connected in the network by bounded-bandwidth links which we refer to as edges. The network is modeled by a graph $G = (V, E)$ where V is the set of processors and E is the set of edges or links between these processors. Here $|V| = n$ and $|E| = m$.

Each node is assigned a unique ID between 1 and $\text{poly}(n)$ and has infinite computational power. Each node has limited topological knowledge and only knows about its incident edges. For the weighted APSP problem we consider, each edge has a positive integer weight bounded by $\text{poly}(n)$. Also if the edges are directed, the corresponding communication channels are bidirectional and hence

the communication network can be represented by the underlying undirected graph U_G of G (this model is also used in [10, 13, 22]).

The computation proceeds in rounds. In each round each processor can send a message of size $O(\log n)$ along edges incident to it, and it receives the messages sent to it in the previous round. The model allows a node to send different message along different edges though we do not need this feature in our algorithm. The performance of an algorithm in the CONGEST model is measured by its round complexity, which is the worst-case number of rounds of distributed communication. Hence the goal is to minimize the round complexity of an algorithm.

1.2 Prior Work

Unweighted APSP. For APSP in unweighted undirected graphs, $O(n)$ -round algorithms were given independently in [12, 20]. An improved $n + O(D)$ -round algorithm was then given in [17], where D is the diameter of the undirected graph. Although this latter result was claimed only for undirected graphs, the algorithm in [17] is also a correct $O(n)$ -round APSP algorithm for directed unweighted graphs. The message complexity of directed unweighted APSP was reduced to $mn + O(m)$ in a recent algorithm [22] that runs in $\min\{2n, n + O(D)\}$ rounds (where D is now the directed diameter of the graph). A lower bound of $\Omega(n/\log n)$ for the number of rounds needed to compute the diameter of the graph in the CONGEST model is given in [6].

Weighted APSP. While unweighted APSP is well-understood in the CONGEST model much remains to be done in the weighted case. For deterministic algorithms, weighted SSSP for a single source can be computed in n rounds using the classic Bellman-Ford algorithm [2, 5], and this leads to a simple deterministic weighted APSP algorithm that runs in $O(n^2)$ rounds. Nothing better was known for the number of rounds for deterministic weighted APSP until our current results.

Exact Randomized APSP Algorithms. Even with randomization, nothing better than n^2 rounds was known for exact weighted APSP until recently, when Elkin [3] gave a randomized weighted APSP algorithm that runs in $\tilde{O}(n^{5/3})$ rounds and this was further improved to $\tilde{O}(n^{5/4})$ rounds in Huang et al. [13]. Both of these results hold with high probability in n .

Deterministic Approximation Algorithms for APSP. There are deterministic algorithms for approximating weighted all pairs shortest path problem, and these run in $\tilde{O}(n)$ rounds for both directed [16] and undirected graphs [4, 11, 16].

2 OVERVIEW OF THE APSP ALGORITHM

Let $G = (V, E)$ be an edge-weighted graph (directed or undirected) with weight function w and with $|V| = n$ and $|E| = m$. The CONGEST model assumes that every message is of $O(\log n)$ -bit size, which restricts $w(e)$ to be an $O(\log n)$ size integer value. However, outside of this restriction imposed by the CONGEST model, our algorithm works for arbitrary edge-weights (even negative edge-weights as long as there is no negative-weight cycle). Given a path p we will use *weight* or *distance* to denote the sum of the weights of the edges on the path and *length* (or sometimes *hops*) to denote the number of edges on the path. We denote the shortest path distance

from a vertex x to a vertex y in G by $\delta(x, y)$. In the following we will assume that G is directed, but the same algorithm works for undirected graphs as well.

An h -hop SSSP tree for G rooted at a vertex r is a tree of height h where the weight of the path from r to a vertex v in the tree is the shortest path distance in G among all paths that have at most h edges. In the case of multiple paths with the same distance from r to v we assume that v chooses the parent vertex with minimum id as its parent in the h -hop SSSP tree. We denote the hop-length of the path from root r to v in this tree by $h_r(v)$. We will use $h = \sqrt{n} \cdot \log n$ in our algorithm.

Our overall APSP algorithm is given in Algorithm 1. In Step 1 an h -hop SSSP tree and the associated SSSP distances, $\delta_h(r, v)$, are computed at vertex v for each root $r \in V$.

Step 2 computes a *blocker set* Q of $q = O((n \log n)/h)$ nodes for the collection of h -hop SSSP trees constructed in Step 1. This step is described in detail in the next section, where we describe a distributed implementation of King's sequential method [14]. Our method computes the blocker set Q in $O(nh + (n^2 \log n)/h)$ rounds. We now give the definition of a blocker set for a collection of rooted h -hop trees.

Definition 2.1 (Blocker Set [14]). Let H be a collection of rooted h -hop trees in a graph $G = (V, E)$. A set $Q \subseteq V$ is a *blocker set* for H if every root to leaf path of length h in every tree in H contains a vertex in Q . Each vertex in Q is called a *blocker vertex* for H .

In Step 3 of Algorithm 1 we compute $\delta(c, v)$ for each $c \in Q$ and for all $v \in V$. In Step 4 each blocker vertex c broadcasts all of the $\delta_h(v, c)$ values it computed in Step 1. Finally, in Step 5 each node v computes $\delta(u, v)$ for each $u \in V$ using the values it computed or received in the earlier steps. More specifically, v computes $\delta(u, v)$ as:

$$\delta(u, v) = \min \left\{ \delta_h(u, v), \min_{c \in Q} (\delta_h(u, c) + \delta(c, v)) \right\} \quad (1)$$

LEMMA 2.2. *The $\delta(u, v)$ values computed at each v in Step 5 of Algorithm 1 are the correct shortest path distances.*

PROOF. Fix vertices u, v and consider a shortest path p from u to v . If p has at most h edges then $w(p) = \delta_h(u, v)$ and this value is directly computed at v in Step 1. Otherwise by the property of the blocker set Q we know that there is a vertex $c \in Q$ which lies along p within the h -hop SSSP tree rooted at u that is constructed in Step 1. Let p_1 be the portion of p from u to c and let p_2 be the portion from c to v . So $w(p_1) = \delta_h(u, c)$, $w(p_2) = \delta(c, v)$ and $w(p) = w(p_1) + w(p_2)$.

The value $\delta_h(u, c)$ is received by v in the broadcast step for center c in Step 4. The value $\delta(c, v)$ is computed at v when SSSP with root c is computed in Step 3. Hence v has the information needed to compute $\delta(u, v)$ in Step 5 for each u using Equation 1. $\square$

We now bound the number of rounds needed for each step in Algorithm 1 (other than Step 2). For this we first state bounds for some simple primitives that will be used to execute these steps.

LEMMA 2.3. *Given a source $s \in V$, using the Bellman-Ford algorithm:*

Algorithm 1 Main Algorithm

-
- 1: **for each** $v \in V$ in sequence: compute an h -hop SSSP rooted at v by computing at each $x \in V$ the h -hop shortest path distances $\delta_h(v, x)$, the parent pointer in the SSSP tree, and the hop-length $h_v(x)$ in T_v .
 - 2: Compute a blocker set Q of size $\Theta((n \log n)/h)$ for the h -hop SSSP trees computed in Step 1 using Algorithm 2.
 - 3: **for each** $c \in Q$ in sequence: compute SSSP with root c and $\delta(c, v)$ at each $v \in V$
 - 4: **for each** $c \in Q$ in sequence: broadcast to all other nodes c 's ID and the $\delta_h(v, c)$ values it computed for each v in Step 1.
 - 5: **Local step at each node:** At each v compute $\delta(u, v)$ values for each $u \in V$ by using equation 1 with $\delta_h(u, v)$ from Step 1, the $\delta(c, v)$ values from Step 3, and the $\delta_h(u, c)$ values received in Step 4.
-

- (a) the shortest path distance $\delta(s, v)$ can be computed at each $v \in V$ in n rounds.
- (b) the h -hop shortest path distance $\delta_h(s, v)$, the hop length $h_s(v)$, and parent pointer in the h -hop SSSP tree rooted at s can be computed at each $v \in V$ in h rounds.

LEMMA 2.4. A node v can broadcast k local values to all other nodes reachable from it deterministically in $O(n + k)$ rounds.

PROOF. We construct a BFS tree rooted at v in at most n rounds and then we pipeline the broadcast of the k values. The root v sends the i -th value to all its children in round i for $1 \leq i \leq k$. In a general round, each node x that received a value in the previous round sends that value to all its children. It is readily seen that the i -th value reaches all nodes at hop-length d from v in the BFS tree in round $i + d - 1$, and this is the only value that node x receives in this round. $\square$

LEMMA 2.5. All $v \in V$ can broadcast a local value to every other node they can reach in $O(n)$ rounds deterministically.

PROOF. This broadcast can be done in $O(n)$ rounds in many ways, for example by piggy-backing on an $O(n)$ round unweighted APSP algorithm [17, 22] (and also [12, 20] for undirected graphs) where now each message contains the value sent by source s in addition to the current shortest path distance estimate for source s . $\square$

LEMMA 2.6. Algorithm 1 runs in $O(n \cdot h + (n^2/h) \cdot \log n)$ rounds assuming Step 2 can be implemented to run within this bound.

PROOF. Let the size of the blocker set be $q = \frac{n}{h} \cdot \log n$. Using part (b) of Lemma 2.3, Step 1 can be computed in $O(n \cdot h)$ rounds by computing the n h -hop SSSP trees in sequence. Step 3 can be computed in $O(n \cdot q) = O((n^2/h) \cdot \log n)$ rounds by part (a) of Lemma 2.3. Step 4 can be computed in $O(n \cdot q) = O((n^2/h) \cdot \log n)$ rounds by Lemma 2.4 (using $k = n$). Finally, Step 5 involves only local computation and no communication. This establishes the lemma. $\square$

In the next section we give a deterministic algorithm to compute Step 2 in $O(n \cdot h + (n^2/h) \cdot \log n)$ rounds, which leads to our main theorem (by using $h = \sqrt{n \cdot \log n}$).

THEOREM 2.7. Algorithm 1 is a deterministic distributed algorithm for weighted APSP in directed or undirected graphs that runs in $O(n^{3/2} \cdot \sqrt{\log n})$ rounds in the CONGEST model.

3 COMPUTING A BLOCKER SET DETERMINISTICALLY

The simplest method to find a blocker set is to choose the vertices randomly. An early use of this method for path problems in graphs was in Ullman and Yannakakis [23] where a random set of $O(\sqrt{n} \cdot \log n)$ distinguished nodes was picked. It is readily seen that some vertex in this set will intersect any path of $O(\sqrt{n})$ vertices in the graph (and so this set would serve as a blocker set of size $O((n \log n)/h)$ for our algorithm if $h = \sqrt{n}$). Using this observation an improved randomized parallel algorithm (in the PRAM model) was given in [23] to compute the transitive closure. Since then this method of using random sampling to choose a suitable blocker set has been used extensively in parallel and dynamic computation of transitive closure and shortest paths, and more recently, in distributed computation of APSP [13].

It is not clear if the above simple randomized strategy can be derandomized in its full generality. However, for our purposes a blocker set only needs to intersect all paths in the set of hop trees we construct in Step 1 of Algorithm 1. For this, a deterministic sequential algorithm for computing a blocker set was given in King [14] in order to compute fully dynamic APSP. This algorithm computes a blocker set of size $O((n/h) \ln p)$ for a collection F of h -hop trees with a total of p leaves across all trees (and hence p root to leaf paths) in an n -node graph. In our setting $p \leq n^2$ since we have n trees and each tree could have up to n leaves.

King's sequential blocker set algorithm uses the following simple observation: Given a collection of p paths each with exactly h nodes from an underlying set V of n nodes, there must exist a vertex that is contained in at least ph/n paths. The algorithm adds one such vertex v to the blocker set, removes all paths that are covered by this vertex and repeats this process until no path remains in the collection. The number of paths is reduced from p to at most $(1 - h/n) \cdot p$ when the blocker vertex v is removed, hence after $O((n/h) \ln p)$ removals of vertices, all paths are removed. Since p is at most n^2 the size of the blocker set is $O((n \log n)/h)$. King's sequential algorithm for finding a blocker set runs in $O(n^2 \log n)$ deterministic time (a correction to the bound given in King [14]).

We now describe our distributed algorithm to compute a blocker set. As in King [14], for each vertex v in a tree T_x in the collection of trees H we define:

- $score_x(v)$ is the number of leaves at depth h in T_x that are in the subtree rooted at v in T_x ;
- $score(v) = \sum_x score_x(v)$.

Thus, $score(v)$ is the number of root-to-leaf paths of length h in the collection of trees H that contain vertex v . Initially, our distributed algorithm computes all $score_x(v)$ and $score(v)$ for all vertices $v \in V$

Algorithm 2 COMPUTE-BLOCKER

Input: set H of all h -hop trees T_x ; Output: set Q

- 1: **Initialization** [lines 2-6]:
- 2: Run Algorithm 3 to compute scores for all $v \in V$
- 3: For each T_x compute the ancestors of each vertex v in T_x in $Anc_x(v)$ using Algorithm 4
- 4: **for** each $v \in V$ **do**
- 5: **Local Step:** $score(v) \leftarrow \sum_{x \in V} score_x(v)$
- 6: broadcast $score(v)$ to all nodes in V (using Lemma 2.5)
- 7: **Add blocker vertices to blocker set** Q [lines 8-12]:
- 8: **while** there is a node c with $score(c) > 0$ **do**
- 9: **for** each $v \in V$ **do**
- 10: **Local Step:** select the node c with max score as next vertex in Q
- 11: Run Algorithms 5 and 6 to update $score_x(v)$ for each $x \in V$ and $score(v)$
- 12: broadcast $score(v)$ to all nodes in V and receive $score(x)$ from all other nodes x

and all h -hop trees T_x in $O(n \cdot h)$ rounds using Algorithm 3. Then through an all-to-all broadcast of $score(v)$ to all other nodes for all v , all nodes identify the vertex c with maximum score as the next blocker vertex to be removed from the trees and added to the blocker set Q . (In case there are multiple vertices with the maximum score the algorithm chooses the vertex of minimum id having this maximum score. This ensures that all vertices will locally choose the same vertex as the next blocker vertex once they have received the scores of all vertices.) We repeat this process until all scores are zeroed out. By the discussion above (and as observed in [14]) we will identify all the vertices in Q in $O((n \cdot \log n)/h)$ repeats of this process.

What remains is to obtain an $O(n)$ round procedure to update the $score$ and $score_x$ values at all nodes each time a vertex c is removed so that we have the correct values at each node for each tree when the leaves covered by c are removed from the tree.

If a vertex v is a descendant of the removed vertex c in T_x then all paths in T_x that pass through v are removed when c is removed and hence $score_x(v)$ needs to go down to zero for each such tree T_x where v is a descendant of the chosen blocker node c . In order to facilitate an $O(n)$ -round computation of these updated $score_x$ values in each tree at all nodes that are descendants of c , we initially precompute at every node v a list $Anc_x(v)$ all of its ancestors in each tree T_x . This is computed in $O(n \cdot h)$ rounds using Algorithm 4. Thereafter, each time a new blocker vertex c is selected to be removed from the trees and added to Q , it is a local computation at each node v to determine which of the $Anc_x(v)$ sets at v contain c and to zero out $score_x(v)$ for each such x .

The other type of vertices whose scores change after a vertex c is removed are the ancestors of c in each tree. If v is an ancestor of c in T_x then after c is removed $score_x(v)$ needs to be reduced by $score_x(c)$ (i.e., c 's score before it was removed and added to Q) since these paths no longer need to be covered by v . For these ancestor updates we give an $O(n)$ -round algorithm that runs after the addition of each new blocker node to Q and correctly updates the scores for these ancestors in every tree (Algorithm 6). These algorithms together give the overall deterministic algorithm (Algorithm 2) for the computation of the blocker set Q in $O(n \cdot h + (n^2 \log n)/h)$ rounds.

We now give the details of our algorithms. Recall that we use the Bellman-Ford algorithm for Step 1 in Algorithm 1. Hence after that step, for each tree T_s rooted at s every node v in the tree knows its shortest path distance from s , $\delta(s, v)$, its hop length $h_s(v)$ and its parent node in T_s . We also determine for each node its children in T_s . We can compute this in one round for each T_s by have each node send its child status to its parent. Thus after n rounds all nodes know all their children in every tree T_s .

3.1 The Blocker Set Algorithm

Algorithm 2 gives our distributed deterministic method to compute a blocker set. It uses a collection of helper algorithms that are described in the next section. This blocker set algorithm is at the heart of our main algorithm (Algorithm 1, Step 2) for computing the exact weighted APSP.

Step 2 of Algorithm 2 executes Algorithm 3 to compute all the initial scores at all nodes v . Step 3 involves running Algorithm 4 for pre-computing ancestors of each node in every T_x . Step 5 is a local computation (no communication) where all nodes v compute their total score by summing up the scores for all trees T_x to which they belong. And in Step 6, each node v broadcasts its score value to all other nodes.

The while loop in Steps 8-12 of Algorithm 2 runs as long as there is a node with positive score. In Step 10, the node with maximum score is selected as the vertex c to be added to Q (and if there are multiple nodes with the maximum score, then among them the node with the minimum ID is selected, so that the same node is selected locally at every vertex). In Step 11, after blocker vertex c is selected, each node v checks whether it is a descendant of c in each T_x and if so update its score for that tree using Algorithm 5. This is followed by an execution of Algorithm 6 which updates the scores at each node v for each tree T_x in which v is an ancestor of c . Then in Step 12, all the nodes broadcast their score to all other nodes so that they can all select the next vertex to be added to Q . This leads to the following lemma, assuming the results shown in the next section.

LEMMA 3.1. *Algorithm 2 correctly computes the blocker set Q in $O(n \cdot h + n \cdot |Q|)$ rounds.*

Algorithm 3 Compute Initial scores for a node v in T_x

```

1: Initialization [Local Step]: if  $h_x(v) = h$  then  $score_x(v) \leftarrow 1$ 
   else  $score_x(v) \leftarrow 0$ 
2: In round  $r > 0$ :
3: send: if  $r = h - h_x(v) + 1$  then send  $\langle score_x(v) \rangle$  to  $parent_x(v)$ 

4: receive [lines 5-9]:
5: if  $r = h - h_x(v)$  then
6:   let  $I$  be the set of incoming messages to  $v$ 
7:   for each  $M \in I$  do
8:     let  $M = \langle score^- \rangle$  and let the sender be  $w$ 
9:     if  $w$  is a child of  $v$  in  $T_x$  then  $score_x(v) \leftarrow score_x(v) + score^-$ 

```

PROOF. Step 2 runs in $O(n \cdot h)$ rounds (by Lemma 3.2) and so does Step 3 (see Lemma 3.3 below). Step 5 is a local computation and the broadcast in Step 6 runs in $O(n)$ rounds by Lemma 2.5.

The while loop starting in Step 8 runs for $|Q|$ iterations since a new blocker vertex is added to Q in each iteration. In each iteration, Step 10 is a local computation as is the execution of Algorithm 5 in Step 11. Algorithm 6 in Step 11 runs in $O(n)$ rounds (Lemma 3.6). The all-to-all broadcast in Step 12 is the same as the initial all-to-all broadcast in Step 6 and runs in $O(n)$ rounds. Hence each iteration of the while loop runs in $O(n)$ rounds giving the desired bound. $\square$

3.2 Algorithms for Computing and Updating Scores

In this section we give the details of our algorithms for computing initial scores (Algorithm 3) and for updating these scores values once a blocker vertex c is selected and added to the blocker set Q (Algorithms 4-6).

Algorithm 3 gives the procedure for computing the initial scores for a node v in a tree T_x . In Step 1 each leaf node at depth h initializes its score for T_x to 1 and all other nodes set their initial score to 0. In a general round $r > 0$, nodes with $h_x(v) = h + 1 - r$ send out their scores to their parents and nodes with $h_x(v) = h - r$ will receive all the scores from its children in T_x and set its score equal to the sum of these received scores (Steps 5-9).

LEMMA 3.2. *Algorithm 3 computes the initial scores for every node v in T_x in $O(h)$ rounds.*

PROOF. The leaves at depth h correctly initialize their score to 1 locally in Step 1. Since we only consider paths of length h from the root x to a leaf, it is readily seen that a node v that is $h_x(v)$ hops away from x in T_x will receive scores from its children in round $h - h_x(v)$ and thus will have the correct $score_x(v)$ value to send in Step 3. $\square$

For every $x \in V$, every node $v \in T_x$ will run this algorithm to compute their score in T_x . Since every run of Algorithm 3 for a given x takes h rounds, all the initial scores can be computed in $O(n \cdot h)$ rounds.

Algorithm 4 describes our algorithm for precomputing the ancestors of each node v in a tree T_x of height h . In round 1, every node v sends its ID to its children in T_x as described in Step 5. And in a general round r , v sends the ID of the ancestor that it received in

Algorithm 4 ANCESTORS (v, x) : Algorithm for computing ancestors of node v in T_x at round r

```

1: Initialization [Local Step]:  $Anc_x(v) \leftarrow \phi$ 
2: In round  $r > 0$ :
3: send [lines 4-8]:
4: if  $r = 1$  then
5:   send  $\langle v \rangle$  to  $v$ 's children in  $T_x$ 
6: else
7:   let  $\langle y \rangle$  be the message  $v$  received in round  $r - 1$ 
8:   send  $\langle y \rangle$  to  $v$ 's children in  $T_x$ 
9: receive [lines 10-11]:
10: let  $\langle y \rangle$  be the message  $v$  received in this round
11: add  $y$  to  $Anc_x(v)$ 

```

Algorithm 5 Algorithm for updating scores at v when v is a descendant of new blocker node c

Input: blocker vertex c added to Q .

There is no communication in this algorithm, it is entirely a local computation at v .

```

1: if  $score(v) \neq 0$  then
2:   for each  $x \in V$  do
3:     if  $c \in Anc_x(v)$  then
4:        $score(v) \leftarrow score(v) - score_x(v)$ 
5:        $score_x(v) \leftarrow 0$ 

```

round $r - 1$ (Steps 7-8). If a node v receives the ID of an ancestor y , then it immediately adds it to its ancestor set, $Anc_x(v)$ (Steps 10-11).

LEMMA 3.3. *For a tree T_x of height h rooted at vertex x , Algorithm 4 correctly computes the set of ancestors for all nodes v in T_x in $O(h)$ rounds.*

PROOF. We show that all nodes v correctly computes all their ancestors in T_x in the set $Anc_x(v)$ using induction on round r . We show that by round r , every node v has added all its ancestors that are at most r hops away from v .

If $r = 1$, then v 's parent in T_x (say y) would have send out its ID to v in Step 5 and v would have added it to $Anc_x(v)$ in Step 11.

Assume that every node v has already added all ancestors in the set $Anc_x(v)$ that are at most $r - 1$ hops away from v .

Let u be the ancestor of v in T_x that is exactly r hops away from v . Then by induction, $u \in Anc_x(y)$ since u is exactly $r - 1$ hops away from y and thus y must have send u 's ID to v in round r in Step 8 and hence v would have added u to its set $Anc_x(v)$ in round r in Step 11. $\square$

Once we have pre-computed the $Anc_x(v)$ sets for all vertices v and all trees T_x using Algorithm 4, updating the scores at each node for all trees in which it is a descendant of the newly chosen blocker node c becomes a purely local computation. Algorithm 5 describes the algorithm at node v that updates its scores after a vertex c is added as a blocker node to Q . At node v for each given T_x , v checks if $c \in Anc_x(v)$ and if so update its score values in Steps 4-5.

LEMMA 3.4. *Given a blocker vertex c , Algorithm 5 correctly updates the scores of all nodes v such that v is a descendant of c in some tree T_x .*

Algorithm 6 Pipelined Algorithm for updating scores at v for all trees T_x in which v is an ancestor of newly chosen blocker node c

Input: current blocker set Q , newly chosen blocker node c

```

1: Send [lines 2-3]: (only for  $c$ )
2: Local Step at  $c$ : create a list  $list_c$  and for each  $x \in V$  do add an entry  $Z = \langle x, score_x(c) \rangle$  to  $list_c$  if  $score_x(c) \neq 0$ ; then set  $score_x(c)$  to 0 for each  $x \in V$  and set  $score(c)$  to 0
3: Round  $i$ : let  $Z = \langle x, score_x(c) \rangle$  be the  $i$ -th entry in  $list_c$ ; send  $\langle Z \rangle$  to  $c$ 's parent in  $T_x$ 
4: In round  $r > 0$ : (for vertices  $v \in V - Q - \{c\}$ )
5: send [lines 6-8]:
6: if  $v$  received a message in round  $r - 1$  then
7:   let that message be  $\langle Z \rangle = \langle x, score_x(c) \rangle$ .
8:   if  $v \neq x$  then send  $\langle Z \rangle$  to  $v$ 's parent in  $T_x$ 
9: receive [lines 10-11]:
10: if  $v$  receives a message  $M$  of the form  $\langle x, score_x(c) \rangle$  then
11:    $score_x(v) \leftarrow score_x(v) - score_x(c)$ ;  $score(v) \leftarrow score(v) - score_x(c)$ 

```

PROOF. Fix a vertex v and a tree T_x such that v is a descendant of c in T_x . By Lemma 3.3 $c \in Anc_x(v)$, and thus v will correctly update its score values in Steps 4-5. $\square$

We now move to the last remaining part of the blocker set algorithm: our method to correctly update scores at ancestors of the newly chosen blocker node c in each T_x . Recall that if v is an ancestor of c in T_x we need to subtract $score_x(c)$ from $score_x(v)$. Here, in contrast to Algorithms 4 and 5 for nodes that are descendants of c in a tree, we do not precompute anything. Instead we give an $O(n)$ -round method in Algorithm 6 to correctly update scores for each vertex for all trees in which that vertex is an ancestor of c .

Before we describe Algorithm 6 we establish the following lemma, which is key to our $O(n)$ -round method.

LEMMA 3.5. Fix a vertex c . For each root vertex $x \in V - \{c\}$, let $\pi_{x,c}$ be the path from x to c in the h -hop SSSP tree T_x . Let $T = \cup_{x \in V - \{c\}} \{e \mid e \text{ lies on } \pi_{x,c}\}$, i.e., T is the set of edges that lie on some $\pi_{x,c}$. Then T is an in-tree rooted at c .

PROOF. If not, there exists some $x, y \in V - \{c\}$ such that $\pi_{x,c}$ and $\pi_{y,c}$ coincide first at some vertex z and the subpaths in $\pi_{x,c}$ and $\pi_{y,c}$ from z to c are different.

Consider the first vertex z' at which these two paths coincide after diverging from z (such a vertex exists since the two paths meet at c). Let the subpath from z to z' in $\pi_{x,c}$ be $\pi_{z,z'}^1$, and the corresponding subpath in $\pi_{y,c}$ be $\pi_{z,z'}^2$. Similarly let $\pi_{x,z}$ be the subpath of $\pi_{x,c}$ from x to z and let $\pi_{y,z}$ be the subpath of $\pi_{y,c}$ from y to z .

Clearly both $\pi_{z,z'}^1$ and $\pi_{z,z'}^2$ have equal weight (otherwise one of $\pi_{x,c}$ or $\pi_{y,c}$ cannot be a shortest path). Thus the path $\pi_{x,z} \circ \pi_{z,z'}^2$ is also a shortest path.

Let (a, z') be the last edge on the path $\pi_{z,z'}^1$, and (b, z') be the last edge on the path $\pi_{z,z'}^2$. Nodes a and b must be different by our choice of vertex z' . Since the path $\pi_{x,z'}$ has (a, z') as the last edge and we break ties using the IDs of the vertices, hence $ID(a) < ID(b)$. But then the shortest path $\pi_{y,z'}$ must also have chosen (a, z') as the last edge and hence $\pi_{y,z} \circ \pi_{z,z'}^1$ must be the subpath of path $\pi_{y,c}$, resulting in a contradiction. $\square$

Lemma 3.5 allows us to re-cast the task for ancestor nodes to the following (where we use the notation in the statement of Lemma 3.5): the new blocker node c needs to send $score_x(c)$ to all nodes on $\pi_{x,c}$ for each tree T_x . Recall that in the CONGEST model for directed graphs the graph edges are bi-directional. Hence this task can be accomplished by having c send out $score_x(c)$ for each tree T_x (other than T_c) in $n - 1$ rounds, one score per round (in no particular order) along the parent edge for T_x . Each message $\langle x, score_x(c) \rangle$ will move along edges in $\pi_{x,c}$ (in reverse order) along parent edges in T_x from c to x . Consider any node v . In general it will be an ancestor of c in some subset of the $n - 1$ trees T_x . But the characterization in Lemma 3.5 establishes that the incoming edge to v in all of these trees is the same edge (u, v) and this is the unique edge on the path from c to v in tree T (T is defined in the statement of Lemma 3.5). In fact, the messages for all of the trees in which v is an ancestor of c will traverse exactly the same path from c to v . Hence, for the messages sent out by c for the different trees in $n - 1$ different rounds (one for each tree other than T_c), if each vertex simply forwards any message $\langle x, score_x(c) \rangle$ it receives to its parent in tree T_x all messages will be pipelined to all ancestors in $n - 1 + h$ rounds. This is what is done in Algorithm 6, whose steps we describe below, for completeness.

Step 2 of Algorithm 6 is local computation at the new blocker vertex c where for each T_x to which c belongs, c adds an entry $\langle x, score_x(c) \rangle$ to a local list $list_c$. In round i , c sends the i -th entry in its list, say $\langle y, score_y(c) \rangle$, to its parent in T_y . For node v other than c , in a general round $r > 0$, if v receives a message for some $x \in V$ it updates its score value for x (Steps 10-11) and then forwards this message to its parent in T_x in round $r + 1$ (Step 6-8).

LEMMA 3.6. Given a new blocker vertex c , Algorithm 6 correctly updates the scores of all nodes v in every tree T_x in which v is an ancestor of c in $O(n + h)$ rounds.

PROOF. Correctness of Algorithm 6 was argued above. For the number of rounds, c sends out its last message in round $n - 1$, and if $\pi_{v,c}$ has length k then v receives all messages sent to it by round $n - 1 + k$. Since we only have h -hop trees $k \leq h$ for all nodes, and the lemma follows. $\square$

4 CONCLUSION

We have presented a new distributed algorithm for the exact computation of weighted all pairs shortest paths in both directed and undirected graphs. This algorithm runs in $O(n^{3/2} \cdot \sqrt{\log n})$ rounds and is the first $o(n^2)$ -round deterministic algorithm for this problem in the CONGEST model. At the heart of our algorithm is a deterministic algorithm for computing blocker set. Our blocker set construction may have applications in other distributed algorithms that need to identify a relatively small set of vertices that intersect all paths in a set of paths with the same (relatively long) length.

The main open question left by our work is to improve the bound on the number of rounds for deterministic weighted APSP. In the unweighted case APSP can be computed in $O(n)$ rounds [12, 17, 20, 22], and weighted APSP can be computed in $\tilde{O}(n^{5/4})$ rounds [13] w.h.p. with randomization. In this paper we give the first $o(n^2)$ rounds deterministic algorithm for weighted APSP. Also very recently in [1], the first two authors presented a new pipelined algorithm for Step 1 of our Algorithm 1 that runs in $O(n\sqrt{h})$ rounds (improved from $O(nh)$ that we have here). This, together with some changes to the blocker set algorithm in Step 2, gives an $\tilde{O}(n^{4/3})$ rounds deterministic algorithm for APSP in the Congest model. This leaves open the question of whether we can match the current randomized bound with a deterministic algorithm, and whether we can improve the bound for the number of rounds to below $n^{5/4}$ either deterministic or randomized.

REFERENCES

- [1] Udit Agarwal and Vijaya Ramachandran. 2018. A Faster Deterministic Distributed Algorithm for Weighted APSP Through Pipelining. (2018). Manuscript.
- [2] Richard Bellman. 1958. On a routing problem. *Quarterly of applied mathematics* 16, 1 (1958), 87–90.
- [3] Michael Elkin. 2017. Distributed exact shortest paths in sublinear time. In *Proc. STOC*. ACM, 757–770.
- [4] Michael Elkin and Ofer Neiman. 2016. Hopsets with constant hopbound, and applications to approximate shortest paths. In *Proc. FOCS*. IEEE, 128–137.
- [5] Lester R Ford Jr. 1956. *Network flow theory*. Technical Report. RAND CORP SANTA MONICA CA.
- [6] Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. 2012. Networks Cannot Compute Their Diameter in Sublinear Time. In *SODA '12*. 1150–1162.
- [7] Robert G. Gallager, Pierre A. Humblet, and Philip M. Spira. 1983. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Prog. Lang. Syst.* 5, 1 (1983), 66–77.
- [8] Juan A Garay, Shay Kutten, and David Peleg. 1998. A sublinear time distributed algorithm for minimum-weight spanning trees. *SIAM J. Comp.* 27, 1 (1998), 302–316.
- [9] Mohsen Ghaffari. 2015. Near-optimal scheduling of distributed algorithms. In *Proc. PODC*. ACM, 3–12.
- [10] Mohsen Ghaffari and Jason Li. 2017. Improved Distributed Algorithms for Exact Shortest Paths. *arXiv preprint arXiv:1712.09121* (2017).
- [11] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. 2016. A deterministic almost-tight distributed algorithm for approximating single-source shortest paths. In *Proc. STOC*. ACM, 489–498.
- [12] Stephan Holzer and Roger Wattenhofer. 2012. Optimal Distributed All Pairs Shortest Paths and Applications. In *PODC '12*. 355–364.
- [13] Chien-Chung Huang, Danupon Na Nongkai, and Thatchaphol Saranurak. 2017. Distributed Exact Weighted All-Pairs Shortest Paths in $\tilde{O}(n^{5/4})$ Rounds. In *Proc. FOCS*. IEEE, 168–179.
- [14] Valerie King. 1999. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *Proc. FOCS*. IEEE, 81–89.
- [15] Shay Kutten and David Peleg. 1998. Fast distributed construction of small-k dominating sets and applications. *J. Algorithms* 28, 1 (1998), 40–66.
- [16] Christoph Lenzen and Boaz Patt-Shamir. 2015. Fast partial distance estimation and applications. In *Proc. PODC*. ACM, 153–162.
- [17] Christoph Lenzen and David Peleg. 2013. Efficient distributed source detection with limited bandwidth. In *Proc. PODC*. ACM, 375–382.
- [18] Danupon Nanongkai. 2014. Distributed approximation algorithms for weighted shortest paths. In *Proc. STOC*. ACM, 565–573.
- [19] David Peleg. 2000. Distributed computing: A Locality-sensitive Approach. *SIAM Monographs on discrete mathematics and applications* 5 (2000).
- [20] David Peleg, Liam Roditty, and Elad Tal. 2012. Distributed algorithms for network diameter and girth. In *Proc. ICALP*. Springer, 660–672.
- [21] David Peleg and Vitaly Rubinfeld. 1999. A near-tight lower bound on the time complexity of distributed mst construction. In *Proc. FOCS*. IEEE, 253–261.
- [22] Matteo Pontecorvi and Vijaya Ramachandran. 2018. Distributed algorithms for directed betweenness centrality and all pairs shortest paths. *arXiv preprint arXiv:1805.08124* (2018).
- [23] Jeffrey D Ullman and Mihalios Yannakakis. 1991. High-probability parallel transitive-closure algorithms. *SIAM J. Comp.* 20, 1 (1991), 100–125.