

MST Construction in $O(\log \log n)$ Communication Rounds

(Extended Abstract)

Zvi Lotker
Dept. of Electrical Engineering
Tel Aviv University
Tel Aviv 69978
Israel
zvilo@eng.tau.ac.il

Boaz Patt-Shamir^{*}
Cambridge Research Laboratory
Hewlett Packard
One Cambridge Center, Cambridge MA 02142
USA
Boaz.PattShamir@HP.com

Elan Pavlov
Dept. of Computer Science
Hebrew University
Jerusalem 91904
Israel
elan@cs.huji.ac.il

David Peleg[†]
Dept. of Computer Science and Applied Math.
Weizmann Institute of Science
Rehovot 76100
Israel
peleg@wisdom.weizmann.ac.il

ABSTRACT

We consider a simple model for overlay networks, where all n processes are connected to all other processes, and each message contains at most $O(\log n)$ bits. For this model, we present a distributed algorithm that constructs a minimum-weight spanning tree in $O(\log \log n)$ communication rounds, where in each round any process can send a message to each other process. This result is the first to break the $\Omega(\log n)$ parallel time complexity barrier with small message sizes.

1. INTRODUCTION

A minimum-weight spanning tree (MST) is one of the most useful distributed constructs, as it minimizes the cost associated with global operations such as broadcasts and convergecasts. In this paper we present an MST construction algorithm that works in $O(\log \log n)$ communication rounds, where in each round each process can send $O(\log n)$ bits to each other process (more intuitively, each message can contain the identity and weight of only a constant number of edges). Our result shows that MST can be constructed with very little pairwise communication: throughout the execu-

tion of the algorithm, each pair of processes exchanges at most $O(\log n \log \log n)$ bits; the overall number of bits communicated is $\Theta(n^2 \log n)$, which is optimal. The algorithm extends to larger message sizes, in the sense that the number of communication rounds is $O(\log \frac{1}{\epsilon})$ if each message can contain n^ϵ bits for some $\epsilon > 0$. Note that if messages are not restricted in size, then MST can be trivially constructed in a single round of communication: each process sends all its information to all its neighbors, and then each node can locally compute the MST.

The number of communication rounds dominates the time complexity in situations where latency is high and bandwidth is scarce. This may be the situation in some *overlay networks*. Briefly, the idea in overlay networks is to think of the underlying communication network (e.g., the Internet) as a “black box” that provides reliable point-to-point communication. On top of that network run distributed applications. This approach (whose precursor is the Internet’s “end-to-end argument” [13]) is different from classical distributed models, where processes reside in networks nodes (i.e., switches or routers), and thus their implementation would require using low-level communication. Rather, the pragmatic view now is that distributed applications create their own overlay network by choosing which pairs of local processes will communicate directly according to various criteria. The concept of overlay networks is central to areas such as multicast or content distribution networks (see, e.g., [8] and references therein), peer-to-peer systems (for example, Chord [14]), and others.

1.1 Related Work

ST construction is well studied as a sequential optimization problem (see, e.g., [15, 9]). Distributed MST constructions are presented in [6, 3] (and see references in [11]). These classical distributed algorithms are oriented towards mini-

^{*}On leave from the Dept. of Electrical Engineering, Tel Aviv University, Tel Aviv 69978, Israel.

[†]Supported in part by a grant from the Israel Science Foundation.

mizing the total number of messages in general networks, and their time complexity is inherently $\Omega(\log n)$, even when run on fully connected graphs. The model we use in this paper is a special case of the model studied in [7, 12, 10]: in these papers, each message has $O(\log n)$ bits, but the fully connected graph is not directly considered. The best previously known upper bound for fully connected graphs in this model is $O(\log n)$ communication rounds. This bound holds also for graphs of diameter 2 [10]. (It is known that the number of rounds jumps at least to $\Omega(n^{1/4})$ when the diameter of the network is 3 or more [10, 12].)

The parallel time complexity of MST construction depends on the particular architecture considered, but we are not aware of any sub-logarithmic time algorithm that uses small messages. For the PRAM model, there are quite a few $O(\log n)$ algorithms, including a deterministic one for the CRCW model [4], and a randomized one for the EREW model [5]. Adler *et al.* [1] study the total number of bits that must be communicated in the course of MST construction problem under various parallel architectures. For our model, their results imply that the worst-case number of bits that need to be communicated throughout the execution of the algorithm is $\Omega(n^2 \log n)$.

1.2 System Model

In the underlying formal model, the system is represented by a complete n -node weighted undirected graph $G = (V, E, \omega)$ where $\omega(e)$ denotes the weight of edge $e \in E$. Each node has a distinct ID of $O(\log n)$ bits. Each node knows all the edges it is incident to (and hence, since the graph is a clique, each node knows about all other nodes in the system). An execution of the system proceeds in asynchronous steps: in a receive step, a node receives some of the messages sent to it in previous steps. In a send step, a node makes a local computation and sends messages to the other nodes in the system. Each message may be different, and we require that each message contains at most $O(\log n)$ bits. (The results are extended to larger message sizes in Section 4.) We assume that messages may be delayed arbitrarily, but are never lost or corrupted. The time complexity of an algorithm in the asynchronous model is measured by normalizing the scale so that the longest message delivery time is one unit.

Simplification: The Synchronous Model. In the synchronous model, computation advances in global rounds, where in each round processes send messages, receive them, and do some local computation. This model is much more convenient as a programming mode. Fortunately, since we assume that the system is reliable, we may apply a *synchronizer* that allows us to present the algorithm in the synchronous model. Specifically, we use the β synchronizer of Awerbuch [2]. Let us outline the idea briefly. Assume that we have an algorithm SA for the synchronous model. The execution in the asynchronous model is done as follows. A process starts the next round only after receiving a special “proceed” message from a distinguished node v_0 (say, the node with the lowest ID in the system). It then sends messages according to SA . For each SA message received, the receiver node sends an “ack” message back to the sender; when a sender had received acknowledgements to all the messages it sent, the

sender forwards a “safe” message to v_0 ; when v_0 receives “safe” messages from all nodes in the system, it sends a “proceed” messages to all other nodes, which may then send their next round SA messages. Note that since we assume that the graph is fully connected, this transformation incurs only a constant blowup in the message complexity and in time complexity. We shall henceforth use the synchronous model, but we emphasize that the algorithm works for the asynchronous model as described above.

1.3 The MST Construction Problem

We assume that in the initial state, the input to each node $v \in V$ consists of the the weights of all its incident edges $\omega(v, u)$ for all $u \in V \setminus \{v\}$. Edge weights are assumed to be integers that can be represented using $O(\log n)$ bits. Without loss of generality, we assume that all the edge weights are distinct and thus the MST is unique. When our algorithm halts, all nodes know the full list of all $n - 1$ edges in the MST of G .

2. ALGORITHM DESCRIPTION

In this section we describe the algorithm. In Section 2.1 we give an overview of the main ideas. In Section 2.2 we specify the main algorithm, and in Sections 2.3 and 2.4 we specify local subroutines used by the main algorithm.

2.1 Overview

The algorithm operates in phases: Each phase takes $O(1)$ rounds, and there are at most $O(\log \log n)$ phases. At the end of each phase $k \geq 0$, the nodes of G are partitioned into disjoint clusters $\mathcal{F}^k = \{F_1^k, \dots, F_{m_k}^k\}$, $\bigcup_i F_i^k = V$. For each cluster $F \in \mathcal{F}^k$, the algorithm selects also a spanning subtree $T(F)$. The partition \mathcal{F}^k and the corresponding subtree collection $\mathcal{T}^k = \{T(F) \mid F \in \mathcal{F}^k\}$, including the weights of the edges in those subtrees, are known to every vertex in the graph. (For notational consistency, we think of the initial situation at the beginning of Phase 1 as the end of an “imaginary” phase 0, with each node forming a singleton cluster, i.e., $\mathcal{F}^0 = \{F_1^0, \dots, F_n^0\}$ where $F_i^0 = \{v_i\}$ for every $1 \leq i \leq n$.)

Define a *fragment* to be a connected subtree of the MST. For a set of nodes $F \subseteq V$, denote by $\mathcal{T}(F)$ the subgraph of MST induced by F . With these notations, we can state following invariant, satisfied by the algorithm at the end of each phase $k \geq 0$:

For every cluster F , $T(F) = \mathcal{T}(F)$, namely the spanning subtree selected for F is a fragment.

In our model, it is easy for the nodes of each cluster to learn, in constant time, the lightest edge to every other cluster. Hence intuitively, it is possible to “contract” each cluster C into a vertex v_C , thus creating a smaller logical graph \hat{G} , and continue working on this logical graph. (In practice, each real vertex belonging to some cluster C knows the weight of the edge connecting its vertex v_C to every other vertex in \hat{G} . The operations of each vertex v_C of the logical graph \hat{G} are carried out by the real vertices belonging to the cluster C , or by a single representative called the *leader* of C ,

denoted $\ell(C)$.) This enables us to simulate the usual MST construction process for \hat{G} , based on growing fragments of the MST by examining the edges one by one in increasing order of weight and including in the MST each inspected edge that is the minimum-weight outgoing edge (MWOE) of its fragment. This can be done in $O(\log n)$ time.

To reduce the time complexity to $O(\log \log n)$, it is necessary to speed up the process by making the cluster sizes grow quadratically in each phase. The main idea used for achieving this growth rate is the following. Essentially, we would like to provide every vertex v_C in the logical graph \hat{G} with information about additional edges in \hat{G} , beyond its own. In particular, if we were somehow able to let every vertex v_C learn the *entire* topology of \hat{G} , then we could finish the MST construction for \hat{G} in a single step by asking each vertex in the graph to compute the MST locally. Unfortunately, such information exchange seems to require too much time. On the positive side, denoting the minimum cluster size by N , it is possible for the (N or more) members of each cluster to inform a distinguished vertex v_0 of the graph, in constant time, of the N lightest edges connecting their cluster to other clusters, by appropriately sharing the workload of this task among them. (For concreteness, we assume that v_0 is the node with the smallest ID in the system.)

Subsequently, we now face a special sub-task of the MST construction problem to solve in v_0 . This node now has a partial picture of the logical graph \hat{G} , consisting of all the vertices v_C but only some of the edges connecting them, particularly, the N lightest edges emanating from each vertex of \hat{G} (to N other vertices). It is now necessary to perform (locally) as many *legal* “fragment merging” steps as possible on the basis of this information. That is, we would like to sort the edges known to us by increasing order of weight, examine them one by one, and add edges that are the MWOE of one of the two fragments they connect, so long as we can be sure of that fact. So the question becomes, when is it “dangerous” to continue the merging steps, in the absence of information about the weights of the edges unknown to us.

The answer to this question is, that it is perfectly safe to continue merging a fragment F (in the logical graph \hat{G}), so long as for each vertex v_C in F , we have still not inspected *at least one* of its N lightest edges (which is known to us by assumption). However, once we have already inspected all the edges of some vertex v_C in the fragment F , it becomes dangerous to continue attempting to merge the fragment over edges known to us, as it is possible that the true MWOE of F is the $(N+1)$ st lightest edge emanating from v_C , which is not known to us (yet is lighter than any edge emanating from C that we do know of at this moment).

The crucial observation is that this “safety rule” still allows us to grow each of the fragments to contain at least $N+1$ vertices of \hat{G} . This means that the clusters of the next phase will be of minimum size $\Omega(N^2)$.

An interesting observation is that even when we can no longer identify the MWOE of some fragment F , we may still be able to safely merge F with some other fragment F' . This may still be legitimate if we can ascertain that the edge

connecting F and F' is the MWOE of F' .

Finally, after constructing locally the new fragments, v_0 sends out the identity of the edges added to the chosen set. This can be done in constant time by letting v_0 send each edge to a different intermediate node, which will broadcast that edge to all other nodes.

2.2 The main algorithm

In the algorithm, whenever a node is instructed to send a message containing the edge $e = (u, v)$, this should be interpreted as a message including the ID’s of its two endpoints, $ID(u)$ and $ID(v)$, as well as the edge weight $\omega(e)$.

We now describe the steps taken in phase k for all $1 \leq k \leq \log \log n$. Let v_0 denote the node whose ID is minimal among all nodes in the graph.

Phase k : Code for node v in cluster F of size $N = |F|$

Input: A set of chosen edges. The set of connected components defined by this set is the set of clusters \mathcal{F}^{k-1} . For each cluster $F' \in \mathcal{F}^{k-1}$, $\ell(F')$ is the node with the minimal ID in F' .

1. (a) Compute the minimum-weight edge $e(v, F')$ that connects v to (any node of) F' for all clusters $F' \neq F$.
 (b) Send $e(v, F')$ to $\ell(F')$ for all clusters $F' \neq F$.
 2. **If $v = \ell(F)$ then**
 - (a) Using the messages received from Step 1, compute the lightest edge between F' and F , for every other cluster F' .
 - (b) Perform (locally) Procedure **Cheap_Out**. This procedure (described below) does the following:
 - It selects a set $\mathcal{A}(F)$ containing the N cheapest edges that go out of F to $N = |F|$ distinct clusters; and
 - It appoints for each such edge e a *guardian* node $g(e)$ in F , ensuring that each node in F is appointed as guardian to at most one edge.
 3. Let $e' \in \mathcal{A}(F)$ be the edge for which v was appointed as guardian, i.e., such that $g(e') = v$. Send e' to v_0 , the node with the minimal ID in the graph.
 (At the end of this step, v_0 knows all the edges in the set $\mathcal{A} = \bigcup_{F' \in \mathcal{F}^{k-1}} \mathcal{A}(F')$.)
 4. **If $v = v_0$ then**
 - (a) Perform (locally) Procedure **Const_Frags**. This procedure (described below) computes E^k , the new set of edges to add.
 - (b) For each edge $e \in E^k$, send a message to $g(e)$.
 5. **If v receives a message from v_0 that $e \in E^k$, then v sends e to all nodes in the graph.**
 6. Each node adds all edges in E^k and computes \mathcal{F}^k .
-

2.3 Procedure Cheap_Out

The local procedure `Cheap_Out` is invoked by cluster leaders in each phase, and it operates as follows at the leader of cluster F with $|F| = N$ at phase k .

Input: Cheapest edge $e(F, F')$ for every $F' \in \mathcal{F}^{k-1}$.

1. Sort the input edges in increasing order of weight.
 2. Let $\mu = \min\{N, |\mathcal{F}^{k-1}|\}$.
 3. Define $\mathcal{A}(F)$ to be the first μ edges in the sorted list.
 4. Sort the nodes of F by increasing order of ID.
 5. Appoint the i th node of F as the guardian of the i th edge added to $\mathcal{A}(F)$.
 6. For each node $u \in F$: send the edge to which u is appointed.
-

2.4 Procedure Const_Frags

The local procedure `Const_Frags` is invoked only by the distinguished node v_0 , and it operates as follows. It receives as input the initial partition \mathcal{F}^{k-1} , the spanning subtree collection \mathcal{T}^{k-1} and the set of edges for inspection, \mathcal{A} . Its output is a set of edges E^k , which defines a new partition \mathcal{F}^k and its spanning subtrees \mathcal{T}^k : the edge set of \mathcal{T}^k is the union of the set of edges in \mathcal{T}^{k-1} with the set E^k , and \mathcal{F}^k is the set of connected components of \mathcal{T}^k .

The procedure operates in two stages. In the first stage, it contracts the input clusters into vertices, thus creating a *logical graph* \hat{G} , partitions this logical graph into “super-clusters” and constructs a spanning subtree for each such super-cluster. In the second stage, the procedure transforms the super-clusters and spanning subtrees constructed for \hat{G} into clusters and spanning subtrees for the original graph G .

We now continue with a more detailed description of the two stages. The first stage operates as follows. The procedure starts by creating the logical graph $\hat{G} = (\hat{V}, \hat{E})$, where each input cluster is viewed as a vertex, namely, $\hat{V} = \mathcal{F}^{k-1}$. The edge set \hat{E} consists of the logical edges corresponding to the edges of the set \mathcal{A} . Set the logical edge corresponding to $e = (u, w)$ to be $X(e) = (F, F')$ where $v \in F$ and $w \in F'$. Then $\hat{E} = \{X(e) \mid e \in \mathcal{A}\}$. Each logical edge $X(e)$ is assigned the same weight as e .

Then, the procedure constructs a collection $\hat{\mathcal{F}}$ of super-clusters and a corresponding collection $\hat{\mathcal{T}}$ of spanning subtrees on this logical graph. The construction operates as follows. The procedure first initializes the output partition as $\mathcal{F} = \{\{F\} \mid F \in \mathcal{F}^{k-1}\}$, i.e., each vertex of $\hat{V} = \mathcal{F}^{k-1}$ is a separate super-cluster. The output collection of spanning subtrees is initialized to $\hat{\mathcal{T}} = \emptyset$. The procedure then inspects the edges of \hat{E} sequentially, in increasing order of weight. An inspected logical edge $X(e)$ is added to $\hat{\mathcal{T}}$ if it does not close a cycle with edges already in $\hat{\mathcal{T}}$. Whenever an edge $X(e) = (F_1, F_2)$ is added to $\hat{\mathcal{T}}$, the super-clusters \hat{F}_1 and \hat{F}_2 containing F_1 and F_2 respectively are merged into

one super-cluster \hat{F} , setting $\hat{F} = \hat{F}_1 \cup \hat{F}_2$ and eliminating \hat{F}_1 and \hat{F}_2 , and the corresponding spanning subtrees are fused together into a spanning subtree for the new super-cluster \hat{F} , setting $\hat{T}(\hat{F}) = \hat{T}(\hat{F}_1) \cup \hat{T}(\hat{F}_2) \cup \{X(e)\}$.

In each step during this process, whenever a logical edge $X(e) = (F_1, F_2)$ between two super-clusters \hat{F}_1 and \hat{F}_2 such that $F_1 \in \hat{F}_1$ and $F_2 \in \hat{F}_2$ is inspected, the procedure also considers declaring one or two super-clusters *finished*. In particular:

- If the step resulted in a merge operation creating a new super-cluster $\hat{F} = \hat{F}_1 \cup \hat{F}_2$, then the newly constructed super-cluster \hat{F} is declared finished if one of the following conditions hold:
 - e is the heaviest edge in $\mathcal{A}(F_1)$ or in $\mathcal{A}(F_2)$, or
 - either \hat{F}_1 or \hat{F}_2 is finished.
- If the step did not result in a merge between \hat{F}_1 and \hat{F}_2 , then:
 - The super-cluster \hat{F}_1 is declared finished if e is the heaviest edge in $\mathcal{A}(F_1)$.
 - The super-cluster \hat{F}_2 is declared finished if e is the heaviest edge in $\mathcal{A}(F_2)$.

Also, after every edge inspection step, some of the remaining edges become “dangerous” and are removed from the set \mathcal{A} . A remaining logical edge $X(e) = (F_1, F_2)$, $F_1 \in \hat{F}_1$, $F_2 \in \hat{F}_2$, is still “safe” (i.e., not dangerous) if $e \in \mathcal{A}(F_1)$ and the super-cluster \hat{F}_1 is still unfinished, or if $e \in \mathcal{A}(F_2)$ and the super-cluster \hat{F}_2 is still unfinished. Thus after every edge inspection step, the procedure examines every edge and removes each dangerous edge e from the set \mathcal{A} . The procedure also removes the corresponding logical edge $X(e)$ from \hat{E} . The process terminates once all super-clusters are declared finished (which, as can easily be verified, happens concurrently with the set \mathcal{A} becoming empty).

In the second stage, the procedure transforms the super-clusters and spanning subtrees constructed for \hat{G} into ones for the original graph G . Specifically, for every super-cluster $\hat{F} \in \hat{\mathcal{F}}$ of the logical graph \hat{G} , with spanning subtree $\hat{T}(\hat{F})$, the procedure merges the original clusters included in the super-cluster \hat{F} into a cluster F' of G , and creates the corresponding spanning subtree $T(F')$ for this cluster by merging $\hat{T}(\hat{F})$ together with all the spanning subtrees from the collection \mathcal{T}^{k-1} spanning the original clusters included in the super-cluster \hat{F} , i.e., setting

$$T(F') = \{e \mid X(e) \in \hat{T}(\hat{F})\} \cup \bigcup_{F \in \hat{F}} T(F).$$

It then adds the cluster F' to the output cluster collection \mathcal{F}^k and the spanning subtree $T(F')$ for it into \mathcal{T}^k .

3. ANALYSIS

In this section we prove that the algorithm described in Section 2 is correct and analyze its complexity. It is more convenient to start with the complexity analysis.

3.1 Complexity

The following lemma is the key to the complexity analysis. It bounds from below the growth rate of fragments.

Consider phase k of the algorithm. Let \hat{G} be the logical graph constructed by Procedure `Const_Frags`. Let $\hat{\mathcal{F}}$ be the collection of clusters constructed by Procedure `Const_Frags` for \hat{G} .

LEMMA 3.1. *Every super-cluster in $\hat{\mathcal{F}}$ consists of at least $\mu + 1$ logical vertices of \hat{G} .*

Proof: To establish the lemma, we prove a stronger claim as follows: whenever the procedure declares a super-cluster \hat{F} finished, it contains at least $\mu + 1$ logical vertices of \hat{G} . This claim is proved by structural induction on the super-clusters.

There are three base cases. The first is when \hat{F} is declared finished following a merge step $\hat{F} = \hat{F}_1 \cup \hat{F}_2$ where the two merged super-clusters were unfinished. This merge step was based on the inspection of some logical edge $X(e) = (F_1, F_2)$ such that $F_1 \in \hat{F}_1$ and $F_2 \in \hat{F}_2$. By the algorithm, w.l.o.g. e is the heaviest edge in $\mathcal{A}(F_1)$. As the edges are inspected in increasing weight order, all other edges in $\mathcal{A}(F_1)$ have already been inspected. There are μ such edges, $e_{i_1}, \dots, e_{i_\mu}$, leading to *distinct* original clusters $F_{j_1}, \dots, F_{j_\mu}$. Whenever an edge e_{i_l} was inspected, either the super-clusters containing F_1 and F_{j_l} were merged, or e_{i_l} was found to close a cycle, indicating that F_1 and F_{j_l} already belonged to the same super-cluster. Hence the finished super-cluster \hat{F} contains (at least) the $\mu + 1$ original clusters $F_1, F_{j_1}, \dots, F_{j_\mu}$.

The second base case is when \hat{F} is declared finished following the inspection of some logical edge $X(e) = (F, F_2)$ such that $F \in \hat{F}$ and $F_2 \in \hat{F}_2$, which did not result in a merge. This happens since e is the heaviest edge in $\mathcal{A}(F)$. Again, all μ other edges in $\mathcal{A}(F)$ have already been inspected, and by a similar reasoning as above, the finished super-cluster \hat{F} contains (at least) $\mu + 1$ original clusters. The third base case is the dual case where \hat{F} is declared finished following the inspection of some logical edge $X(e) = (F_1, F)$ such that $F_1 \in \hat{F}_1$ and $F \in \hat{F}$, which did not result in a merge. Again this happens since e is the heaviest edge in $\mathcal{A}(F)$, and the claim follows in the same way.

The inductive claim concerns the case where \hat{F} is declared finished following a merge step $\hat{F} = \hat{F}_1 \cup \hat{F}_2$ where one or both of the two merged super-clusters were finished. In this case, the claim follows directly from the inductive hypothesis. ■

LEMMA 3.2. *The algorithm terminates after no more than $\log \log n + 1$ phases.*

Proof: Denote by β_k the minimum size of a cluster $F \in \mathcal{F}^k$. First note that for all $k \geq 0$,

$$\beta_{k+1} \geq \beta_k(\beta_k + 1) \quad (1)$$

Equation (1) is true by Lemma 3.1, which implies that clusters generated in phase $k + 1$ consist of the union of at least $\beta_k + 1$ clusters of phase k , each containing at least β_k nodes. It follows from Equation (1) that for every $k \geq 1$, $\beta_k \geq 2^{2^{k-1}}$. The lemma follows. ■

The following statement is immediate from the code of the algorithm.

LEMMA 3.3. *Each phase requires $O(1)$ rounds.* ■

We now conclude with the following result.

THEOREM 3.4. *The time complexity of the algorithm is $O(\log \log n)$ rounds, and the overall number of bits communicated is $O(n^2 \log n)$.*

Proof: The time complexity bound follows directly from Lemma 3.2 and Lemma 3.3. For the total number of bits communicated, we account for each step separately as follows. Step 1 of the algorithm involves sending at most $n/2^{2^k-1}$ messages by each node in each phase, and each message contains at most $c \log n$ bits for some constant c . Hence, the number of bits sent over all phases in Step 1 is at most

$$\begin{aligned} \sum_{k=0}^{\log \log n + 1} n \cdot \frac{n \cdot c \log n}{2^{2^k-1}} &= n^2 c \log n \sum_{k=0}^{\log \log n + 1} 2^{-2^k+1} \\ &= O(n^2 \log n). \end{aligned}$$

No messages are sent in Step 2. The number of messages sent in Step 3 of the algorithm in each phase is $O(n)$ over all nodes (since each node receives at most one message), for a total of $O(n \log n \log \log n)$ bits throughout the execution. To account for the number of messages sent in Steps 4 and 5, we bound the total number of messages sent in that step over all nodes and over all phases: note that each edge added to the MST contributes $O(n \log n)$ bits sent at Steps 4 and 5, and since exactly $n - 1$ edges are added to the MST overall, the total number of bits sent in these steps throughout the execution of the algorithm is $O(n^2 \log n)$. The result follows. ■

We note that by the results of Adler *et al.* [1] applied to our model, the minimal number of bits required to solve the MST problem is $\Omega(n^2 \log n)$ in the worst case.

3.2 Correctness

The correctness of the algorithm is proved by the following invariant.

LEMMA 3.5. *In each phase k , for every cluster $F \in \mathcal{F}^k$ constructed by Procedure `Const_Frags`, the corresponding spanning tree is a fragment, namely, $T(F) = T(F)$.*

Proof: By induction on k . The initial partition, \mathcal{F}^0 , trivially satisfies the claim. Now suppose that the collection

\mathcal{T}^{k-1} consists of only ST edges, and consider the collection \mathcal{T}^k constructed in phase k . The spanning subtrees in this collection are composed of spanning subtrees from \mathcal{T}^{k-1} fused together by new edges added by Procedure `Const_Frags`. It suffices to show that every edge added to the trees of \mathcal{T}^k in phase k is indeed an MST edge. For this, we rely on the standard MST construction rule which says that if e is the lightest outgoing edge incident on a fragment, then it belongs to the MST. Consequently, we have to show that whenever Procedure `Const_Frags` selects a logical edge $X(e) = (F_1, F_2)$, $F_1 \in \hat{F}_1$, $F_2 \in \hat{F}_2$, and uses it to merge the two super-clusters \hat{F}_1 and \hat{F}_2 in \hat{G} , then e is the lightest edge outgoing from the corresponding clusters $H_1 = \bigcup_{F \in \hat{F}_1} F$ and $H_2 = \bigcup_{F \in \hat{F}_2} F$ in G .

As the edge e has not been erased prior to this step, necessarily either $e \in \mathcal{A}(F_1)$ and \hat{F}_1 is unfinished, or $e \in \mathcal{A}(F_2)$ and \hat{F}_2 is unfinished. Without loss of generality suppose the former. We claim that in this case e is the lightest outgoing edge incident on H_1 .

Consider some other outgoing edge e' incident on H_1 , and in particular, on a fragment $F_1 \in \hat{F}_1$. Suppose, towards contradiction, that $\omega(e') < \omega(e)$. If e' is included in $\mathcal{A}(F_1)$, then it should have been considered by the procedure before e , and subsequently either added to the spanning subtree $\hat{T}(\hat{F}_1)$ or discarded as an internal edge, in either case contradicting our assumption that e' is an outgoing edge of H_1 (hence $X(e')$ is an outgoing edge of \hat{F}_1). Hence $e' \notin \mathcal{A}(F_1)$. Let $X(e') = (F_1, F_3)$. There may be two reasons why e' was not added to $\mathcal{A}(F_1)$. The first is that some other edge e'' with $X(e'') = (F_1, F_3)$ was already included in $\mathcal{A}(F_1)$ before e' . In that case, $\omega(e'') < \omega(e')$, hence also $\omega(e'') < \omega(e)$. This implies that e'' has already been inspected by the procedure at some earlier step. But then, the clusters F_1 and F_3 must already belong to the super-cluster \hat{F}_1 , hence in \hat{F}_1 , the edge e' is internal, contradiction. The other reason why e' was not added to $\mathcal{A}(F_1)$ is that μ other edges lighter than it were added to $\mathcal{A}(F_1)$ before it. Letting e'' be the heaviest edge in $\mathcal{A}(F_1)$, it follows that $\omega(e'') < \omega(e')$. Hence also $\omega(e'') < \omega(e)$. This means that e'' has already been inspected by the procedure at some earlier step. But then, the cluster \hat{F}_3 that contained it at the end of that step should have been declared finished, upon inspection of its heaviest edge. This would necessitate that \hat{F}_1 is finished now, contradiction. ■

THEOREM 3.6. *The tree produced by the algorithm is a minimum weight spanning tree of the graph.*

Proof: Follows from Lemma 3.5 and the fact that by Lemma 3.2, \mathcal{F}^k contains exactly one cluster for $k > \log \log n$.

4. EXTENSION TO LARGER MESSAGES

In this section we extend the algorithm to a model in which each message can contain any number of bits (so long as it is at least $\log n$). Specifically, we assume that each message may contain $\ell \log n$ bits. The extension of the algorithm to this case is straightforward. It turns out that the asymptotic worst-case number of rounds drops to a constant when the

message size is n^ϵ for $\epsilon > 0$, but $\Theta(\log \log n)$ rounds are required by our algorithm for any polylog message size.

First, we explain how to modify the algorithm to use messages that can contain ℓ edges. The idea is to change Steps 2b (which is the invocation of Procedure `Cheap_Out`) and 3 in the main algorithm so that each node can be the guardian of ℓ edges. Specifically, the modified algorithm is identical to the algorithm of Section 2 except for the following steps.

2b*. Perform (locally) Procedure `Cheap_Out*`. This procedure (described below) does the following:

- It selects a set $\mathcal{A}(F)$ containing the $\ell \cdot N$ cheapest edges that go out of F to $\ell \cdot N$ distinct clusters; and
- It appoints for each such edge e a *guardian* node $g(e)$ in F , ensuring that each node in F is appointed as guardian to at most ℓ edges.

3*. Let $\{e'_1, \dots, e'_\ell\} \subseteq \mathcal{A}(F)$ be the edges for which v was appointed as guardian, i.e., all edges e'_i such that $g(e'_i) = v$. Send $\{e'_1, \dots, e'_\ell\}$ to v_0 , the node with the minimal ID in the graph.
(At the end of this step, v_0 knows all the edges in the set $\mathcal{A} = \bigcup_{F' \in \mathcal{F}^{k-1}} \mathcal{A}(F')$.)

The modified `Cheap_Out*` procedure is identical to Procedure `Cheap_Out` except for the following two steps:

2*. Let $\mu = \min\{\ell \cdot N, |\mathcal{F}^{k-1}|\}$.

5*. Appoint the i th node of F as the guardian of the j th edge added to $\mathcal{A}(F)$ if $j \bmod (\ell \cdot N) = i$.

The correctness of the modification is obvious, as Lemma 3.5 is stated in terms of a general μ , and it relies only on the assumption that $\mathcal{A}(F)$ contains the μ lightest edges connecting F to μ distinct clusters.

The complexity analysis of the generalized algorithm requires a little work. First, we observe that Lemma 3.1 holds without change: it is also stated in terms of a general μ . Lemma 3.3 also holds by assumption that each message can contain ℓ edges, and since each node is the guardian of at most ℓ messages by the modified procedure `Cheap_Out*`. However, Lemma 3.2 holds only for $\ell = \Theta(1)$. Below we analyze two other special classes of values of ℓ .

First, we give a constant upper bound on the number of communication rounds for polynomial size messages.

LEMMA 4.1. *If $\ell \geq n^\epsilon$ for some $\epsilon > 0$, then the worst-case number of rounds of the extended algorithm is $O(\log \frac{1}{\epsilon})$, and the total number of bits communicated is $\Theta(n^2 \log n)$.*

Proof: Let β_k be the smallest possible cluster size after the k th round. By definition, $\beta_0 = 1$. If each guardian node sends ℓ edges, then each cluster merges with at least

$\ell\beta_k$ other clusters in the k th phase. It follows that $\beta_{k+1} \geq (\ell\beta_k + 1)\beta_k \geq \ell\beta_k^2$, and hence $\beta_k > \ell^{2^k}$. Therefore, for $\tau = \log \frac{1}{\epsilon}$ we get

$$\beta_\tau > \ell^{2^{\log \frac{1}{\epsilon}}} = (n^\epsilon)^{\frac{1}{\epsilon}} = n,$$

which means, using Lemma 3.3, that the modified algorithm terminates in $O(\log \frac{1}{\epsilon})$ communication rounds. For the total number of bits communicated by the algorithm, we observe that the only difference in the accounting is that the number of messages sent in Step 3* in each phase, over all nodes, is now $O(n^{1+\epsilon})$, which does not change the asymptotic complexity. ■

Note that Theorem 3.4 can be viewed as a special case of Lemma 4.1, with $\epsilon = \frac{1}{\log n}$.

Next, we show that the worst-case round complexity of the extended algorithm is double-logarithmic even if the messages can have slightly super-polylog size.

LEMMA 4.2. *If $\ell \leq (\log n)^{\frac{(\log n)^c}{\log \log n}} - 1$ for some constant $0 < c < 1$, then the worst-case number of phases of the extended algorithm is larger than $(1 - c) \log \log n$.*

Proof: Using the notation above, we have that $\beta_{k+1} = (\ell\beta_k + 1)\beta_k \leq (\ell + 1)\beta_k^2$, and thus, in the worst case we may have $\beta_k < (\ell + 1)^{2^k}$. Now, suppose that

$$\begin{aligned} \ell &\leq (\log n)^{\frac{(\log n)^c}{\log \log n}} - 1, \\ \tau &\leq (1 - c) \log \log n. \end{aligned}$$

For these values we have

$$\begin{aligned} \beta_\tau &< (\ell + 1)^{2^{(1-c) \log \log n}} \\ &\leq \left((\log n)^{\frac{(\log n)^c}{\log \log n}} \right)^{(\log n)^{1-c}} \\ &= \left(2^{(\log n)^c} \right)^{(\log n)^{1-c}} = 2^{\log n} = n. \end{aligned}$$

Since the extended algorithm may terminate in τ phases only if $\beta_\tau \geq n$, the above inequality shows that in the worst case, the algorithm does not terminate in less than $(1 - c) \log \log n$ phases using the given message size. ■

COROLLARY 4.3. *The number of communication rounds of the extended algorithm is $\Theta(\log \log n)$ for polylogarithmic message sizes.*

5. CONCLUSION

This paper shows that MST can be constructed in sub-logarithmic time even if each message can contain only a constant number of edges. We believe that the algorithm may be useful in some overlay networks. An obvious question we leave open is whether the algorithm can be improved, or is there an inherent lower bound of $\Omega(\log \log n)$ on the number of communication rounds required to construct an MST in this model.

6. REFERENCES

- [1] M. Adler, W. Dittrich, B. Juurlink, M. Kutylowski, and I. Rieping. Communication-optimal parallel minimum spanning tree algorithms. In *Proc. of the 1998 ACM Symposium on Parallel Algorithms and Architecture*, pages 27–36, 1998.
- [2] B. Awerbuch. Complexity of network synchronization. *J. ACM*, 32(4):804–823, Oct. 1985.
- [3] B. Awerbuch. Optimal distributed algorithms for minimum weight spanning tree, counting, leader election and related problems. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 230–240, May 1987.
- [4] B. Awerbuch and Y. Shiloach. New connectivity and MSF algorithms for the shuffle-exchange network and PRAM. *IEEE Trans. Computers*, C-36:1258–1263, 1987.
- [5] R. Cole, P. N. Klein, and R. E. Trajan. Finding minimum spanning forests in logarithmic time and linear work using random sampling. In *Proc. of the 1996 ACM Symposium on Parallel Algorithms and Architecture*, pages 243–250, 1996.
- [6] R. G. Gallager, P. A. Humblet, and P. M. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Prog. Lang. and Syst.*, 5(1):66–77, Jan. 1983.
- [7] J. Garay, S. Kutten, and D. Peleg. A sub-linear time distributed algorithm for minimum-weight spanning trees. *SIAM J. Comput.*, 27:302–316, 1998.
- [8] J. Jannotti, D. Gifford, K. L. Johnson, M. F. Kaashoek, and J. J. W. O’Toole. Overcast: Reliable multicasting with an overlay network. In *4th USENIX OSDI*, pages 197–212, October 2000.
- [9] D. R. Karger, P. N. Klein, and R. E. Tarjan. A randomized linear time algorithm to find minimum spanning trees. *J. ACM*, 42(4):321–328, 1995.
- [10] Z. Lotker, B. Patt-Shamir, and D. Peleg. Distributed MST for constant diameter graphs. In *Proc. 20th Ann. ACM Symp. on Principles of Distributed Computing*, pages 63–71, 2001.
- [11] N. Lynch. *Distributed Algorithms*. Morgan Kaufmann, San Mateo, CA, 1995.
- [12] D. Peleg and V. Rubinovich. Near-tight lower bound on the time complexity of distributed MST construction. *SIAM J. Comput.*, 30:1427–1442, 2000.
- [13] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4):277–288, Nov. 1984.
- [14] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the ACM SIGCOMM ’01 Conference*, San Diego, California, August 2001.
- [15] R. E. Tarjan. *Data Structures and network Algorithms*, chapter 6. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1983.