

Optimal Gossip with Direct Addressing

Bernhard Haeupler
Microsoft Research
1065 La Avenida, Mountain View
Mountain View, CA 94043
haeupler@cs.cmu.edu

Dahlia Malkhi
Microsoft Research
1065 La Avenida, Mountain View
Mountain View, CA 94043
dalia@microsoft.com

ABSTRACT

Gossip algorithms spread information in distributed networks by nodes repeatedly forwarding information to a few random contacts. By their very nature, gossip algorithms tend to be *distributed* and *fault tolerant*. If done right, they can also be *fast* and *message-efficient*. A common model for gossip communication is the random phone call model, in which in each synchronous round each node can PUSH or PULL information to or from a random other node. For example, Karp et al. [FOCS 2000] gave algorithms in this model that spread a message to all nodes in $\Theta(\log n)$ rounds while sending only $O(\log \log n)$ messages per node on average. They also showed that at least $\Theta(\log n)$ rounds are necessary in this model and that algorithms achieving this round-complexity need to send $\omega(1)$ messages per node on average.

Recently, Avin and Elsässer [DISC 2013], studied the random phone call model with the natural and commonly used assumption of *direct addressing*. Direct addressing allows nodes to directly contact nodes whose ID (e.g., IP address) was learned before. They show that in this setting, one can “break the $\log n$ barrier” and achieve a gossip algorithm running in $O(\sqrt{\log n})$ rounds, albeit while using $O(\sqrt{\log n})$ messages per node.

In this paper we study the same model and give a simple gossip algorithm which spreads a message in only $O(\log \log n)$ rounds. We furthermore prove a matching $\Omega(\log \log n)$ lower bound which shows that this running time is best possible. In particular we show that any gossip algorithm takes with high probability at least $0.99 \log \log n$ rounds to terminate. Lastly, our algorithm can be tweaked to send only $O(1)$ messages per node on average with only $O(\log n)$ bits per message. Our algorithm therefore simultaneously achieves the optimal round-, message-, and bit-complexity for this setting. As all prior gossip algorithms, our algorithm is also robust against failures. In particular, if in the beginning an oblivious adversary fails any F nodes our algorithm still, with high probability, informs all but $o(F)$ surviving nodes.

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 the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODC'14, July 15–18, 2014, Paris, France.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2944-6/14/07 ...\$15.00.

<http://dx.doi.org/10.1145/2611462.2611489>.

Categories and Subject Descriptors

F.2.2 [Theory of Computation]: ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY—*Nonnumerical Algorithms and Problems*; F.1.2 [Theory of Computation]: COMPUTATION BY ABSTRACT DEVICES—*Modes of Computation* Parallelism and concurrency

General Terms

Algorithms, Theory, Performance, Reliability

Keywords

direct addressing, gossip, information dissemination, peer-to-peer (P2P), pointer jumping, rumor spreading

1. INTRODUCTION

Gossip algorithms, also called rumor spreading protocols, spread information in a network by having nodes repeatedly forward (limited amounts of) information to a few, often randomly selected, contacts. They have found countless applications in solving coordination and information dissemination tasks in large distributed systems or networks. Their advantages in this setting over alternative approaches is manifold: They are naturally distributed and due to their randomized nature often very *fault tolerant* with respect to many kinds of failures, such as, permanent or transient link-failures or node crashes. If done right, they can also be very efficient in their round-, message-, and bit-complexity.

Two models for gossip communication that have been considered and intensely studied in the literature are the *random phone call model* and *direct addressing*.

The *random phone call model* assumes a complete network and synchronous rounds in which each node is able to PUSH or PULL information to or from a random other node. Both variants spread a message in $O(\log n)$ rounds with $O(\log n)$ messages sent per node [12] if every node communicates in every round. While the logarithmic running time is best possible for this setting, Karp et al. showed in an influential paper [10], that the message-complexity could be reduced to $O(\log \log n)$ message per node on average. They also proved that any address-oblivious algorithm with logarithmic running time needs to send at least $\Omega(\log \log n)$ messages per node on average. Here *address-oblivious* means that the decision of if and what a node sends is independent of the ID of its communication partner(s) this round (see also [1]). Lastly, they demonstrated the *fault-tolerance* of their gossip algorithm. In particular, they remarked that any oblivious F

node failures still lead to all but $O(F)$ nodes being informed in the end, with high probability. Since then the random phone call model has been generalized in many directions, most notably to non-complete network topologies. For example, in [4–6] arbitrary topologies are assumed and the stopping time of the uniform PUSH-PULL gossip is studied where in each round each node initializes a communication with a uniformly random neighbor.

Direct addressing is another natural assumption that has been intensely studied. Direct addressing assumes that nodes have *unique IDs* (e.g., IP addresses) and that in each synchronous round a node can contact any node (or neighbor) whose ID (or IP address) is known to it. Since nodes do not know the topology and are only able to initiate one communication at a time, many algorithms for this setting let nodes contact a random node known to them. For example, in the famous Name-Dropper algorithm of Harchol-Balter et al. [9], nodes repeatedly forward all IDs they know to a random node known to them. They show that this simple algorithm terminates in $O(\log^2 n)$ rounds (using $O(\log^2 n)$ messages per node) if started from any (weakly) connected initial topology. Recently, [11] gave an improved deterministic algorithm which achieves $O(\log n)$ rounds for the same setting. A similar setting but restricting a message to only contain one contact was considered in [8]. Lastly, direct addressing has also been extensively studied in general topologies. For example in [2, 3, 7] it is assumed that all possible communications are given by an arbitrary unknown topology in which nodes only know their neighbors.

Recently Avin and Elsässer [1] showed that adding the natural direct addressing assumption in the random phone call model allows for gossip algorithms that “break the $\log n$ barrier”. Their algorithm runs in $O(\sqrt{\log n})$ rounds however at the cost of having a larger message-complexity of $\Theta(\sqrt{\log n})$ messages per node and also a larger bit-complexity of $O(n \log^{3/2} n + nb \log \log n)$ (compared to the $O(nb \log \log n)$ bit-complexity of [10]). Their algorithm is fault-tolerant in the same way as discussed for [10] above:

THEOREM 1 (THEOREM 1 OF [1]).

In the random phone call model with direct addressing there is a randomized fault-tolerant gossip algorithm that, with high probability, spreads a message of $b = \Omega(\log n)$ bits in $O(\sqrt{\log n})$ rounds using $O(\sqrt{\log n})$ messages per node and a total bit-complexity of $O(n \log^{3/2} n + nb \log \log n)$ bits.

1.1 Our Results

In this paper we study the same setting and improve upon the state of the art by providing a drastically faster algorithm with simultaneously improved message- and bit-complexity:

THEOREM 2.

In the random phone call model with direct addressing there is a randomized fault-tolerant gossip algorithm that, with high probability, spreads a message of $b = \Omega(\log n)$ bits in $O(\log \log n)$ rounds using $O(1)$ messages per node on average and a total bit-complexity of $O(nb)$ bits.

In addition to the faster running time, it is easy to see that the $O(1)$ message-complexity per node and the total $O(nb)$ bit-complexity are optimal. In contrast to [1], which performed worse than [10] on these complexity measures, our

algorithm actually improves on the bounds of [10] and even breaks both message-complexity lower bounds given in [10] for the random phone call model without direct addressing. Lastly, we also prove a matching $\Omega(\log \log n)$ lower bound on the running time of any gossip algorithm in the random phone call model with direct addressing. This shows that our gossip algorithm achieves the optimal round-, message-, and bit-complexity simultaneously.

Our lower bound is particularly strong in the following sense: Instead of showing that any algorithm that takes $c \log \log n$ rounds for a sufficiently small constant c fails with non-negligible probability we prove that any algorithm that takes $0.99 \log \log n$ rounds or less fails with high probability. Our lower bound furthermore holds irrespectively of many model assumptions and even applies to non-gossip algorithms which can contact more than one known node per round:

THEOREM 3. *In the random phone call model with direct addressing any (randomized) gossip algorithm fails to spread a message to all nodes with high probability if less than $0.99 \log \log n$ rounds are used. This remains true if nodes are allowed to use arbitrarily large messages, communicate non address-obliviously and even contact an unbounded number of known nodes per round instead of just one.*

As a last contribution we consider bounds on the number of communications done by a node per round. While in a gossip algorithm each node initializes only one communication per round, nodes might have to respond to multiple requests. In fact, the number of simultaneous requests Δ can be as large as $n - 1$. While this asymmetry is crucially used in essentially all algorithms that use direct addressing, it is unfortunately a questionable assumption for at least some systems and applications (see Section 7 for a more detailed discussion). In this regard, we discuss the easy to prove fact that a sub-logarithmic running time inherently requires $\Delta = \omega(1)$ and show that the best running time possible for any given Δ is $\log n / \log \Delta$. We then give a general framework for working efficiently in a setting with bounded Δ . In particular, we show how our gossip algorithm can efficiently compute a Δ -clustering which can then be used to achieve any point on this trade-off curve between running time and round complexity (subject only to our $\Omega(\log \log n)$ lower bound from Theorem 3). This shows, for example, that one can still have gossip algorithms with an optimal $\Theta(\log \log n)$ running-time while keeping Δ smaller than any polynomial in n , that is, $\Delta = n^{o(1)}$. We believe this to be a useful setting of parameters which seems realistic in many settings.

THEOREM 4. *For any $\Delta = \log^{\omega(1)} n$ there is a randomized gossip algorithm that, with high probability, computes a Δ -clustering in $O(\log \log n)$ rounds using $O(1)$ -messages per node on average but having no node respond to more than Δ requests per round. Given such a Δ -clustering a message broadcast can then be performed easily with an optimal round-, message-, and bit-complexity.*

Organization.

The rest of the paper is organized as follows: In Section 3 we introduce *clusterings*, which are the main abstraction

used to design our gossip algorithms. In Section 4 we give the simplest version of our algorithm which demonstrates the main ideas needed to achieve the $\Theta(\log \log n)$ round complexity. In Section 5 we then explain how to reduce the number of messages sent to $O(1)$ per node on average and achieve an optimal bit-complexity which proves our main theorem. In Section 6 we prove Theorem 3, i.e., our $\Omega(\log \log n)$ round lower bound. In Section 7 we provide our results on gossip algorithms with a bounded amount of communications per node. We defer the discussion of the fault-tolerance of our algorithms to the full version.

2. THE RANDOM PHONE CALL MODEL AND DIRECT ADDRESSING

In this section we give the formal definitions for the random phone call model and the direct addressing assumption used throughout this paper.

The Network.

The random phone call model assumes a complete network consisting of n nodes. Each node has a unique *address* or *ID* consisting of $O(\log n)$ bits, that is, we assume a polynomially large ID space. Initially nodes know their own ID but not the ID of any other node in the system. For convenience we assume that nodes know the number of nodes n . This assumption is without loss of generality since for all problems considered in this paper it is easy to test with high probability whether the algorithm succeeded. This allows for determining the parameter n using the classical guess-test-and-double strategy without increasing the running times by more than a constant factor.

Communication.

Communication takes place in synchronous rounds. In each round a node can choose to initialize one communication. In particular, it can either choose to contact a random node or use direct addressing to contact a node whose ID it knows. Each such contact from a node u to a node v consists of either node u PUSHing a message to v or of node u PULLing a message from v . A message can in principle contain any number of bits but all¹ algorithms throughout this paper only use messages of (minimal) size containing $O(\log n)$ bits. In particular, every message carries either the information to be broadcast, a node count, or $O(1)$ node IDs. Following [1] (and [10]) we say a gossip algorithm is *address-oblivious* if the decision whether to send or not and if so what to send (or respond with) is independent of the communication partner(s) in this round. In particular this means that for every node u the messages PUSHed by u does not depend on the ID of the communication partner and all responses u gives to a PULL are the same in one round. All algorithms presented in this paper are address-oblivious.

Round-, Message-, and Bit-Complexity.

The important figures of efficiency for such a gossip algorithm are the round-complexity, that is, the number of synchronous rounds needed, the message-complexity, that is,

¹The only exceptions are when we share a b -bit message (via a ClusterShare) and for resizing clusters in the Algorithm CLUSTER1, which demonstrates the ideas for achieving the optimal $\Theta(\log \log n)$ round complexity while ignoring message counts and message sizes for simplicity.

the number of messages sent per node on average throughout the execution, and the bit-complexity, that is, the total number of bits communicated in all messages throughout the execution. We only consider and present randomized protocols which complete their task with high probability in the round-, bit-, and message-complexity stated, that is, with probability at least $1 - \frac{1}{n^C}$ for some large chosen constant $C > 1$.

The (Fault-Tolerant) Broadcast Task.

Gossip algorithms can be used for a multitude of different communication and coordination tasks. The algorithms in this paper compute a cluster containing all nodes or, in the case of Theorem 4, a Δ -clustering which can then be used to perform any of these tasks easily and efficiently. To be consistent with prior work we focus on one specific task, namely, broadcasting. In this task there is a message, often also called rumor, consisting of b bits with $b = \Omega(\log n)$ which is initially known to one node (or multiple nodes). The goal is to inform all nodes in the network. In the full version we also address the setting of node-failures. In this setting an oblivious adversary fails any F nodes at the beginning of the execution. Follow [10] and [1], the desired guarantee in this setting is that all but $o(F)$ surviving nodes are clustered/informed. We call an algorithm achieving this *fault-tolerant*.

3. CLUSTERINGS

In this section we introduce clusterings, the main tool used in this paper to achieve efficient gossip.

First, in Section 3.1, we define the concept and implementation of a clustering. In Section 3.2 we then explain how these clustering support efficient implementations of several primitives allowing for coordination and communication between clusters. These primitives are extensively used throughout this paper to implement our fast and message efficient gossip communication algorithms.

3.1 Clustering Definition and Implementation

A *clustering* is a partition of all nodes $v \in V$ in a network into disjoint *clusters* C_1, C_2, \dots, C_k and the set $U = V \setminus \bigcup_i C_i$ of *unclustered nodes*. Each cluster C has one *cluster leader* $l_C \in C$ known to all nodes in C . We call all clustered nodes that are not leaders *cluster followers*. We define the *ID of a cluster* to be the ID of its leader and the *size of a cluster* C to be the number of nodes $|C|$ it contains.

Clusterings are implemented by each node $v \in V$ containing a variable *follow_v*. For any clustered node $v \in C$ this variable contains the ID of the cluster leader $ID(l_C)$ while for all unclustered nodes $v' \in U$ it is set to ∞ . This also allows nodes to determine whether they are a cluster leader or follower by simply comparing their own ID to their *follow* variable.

3.2 Cluster Coordination Primitives

Having a central leader which is known to all nodes in a cluster allows a cluster to coordinate using only a constant number of communication rounds. In one communication round, followers may PUSH information to the leader and thus have the leader collect inputs from the entire cluster. In a second round, followers may PULL a response from the leader. Below, we omit the details concerning PUSHing and PULLing messages, and briefly refer to exchanges such as

above as followers sending information (if any) to the leader and pulling one directive from it.

Following is a list of cluster-macros that make use of the power of clusters, and will be used later in our constructions. All these cluster primitives require only a constant number of rounds and therefore also a constant number of messages per node.

- **ClusterActivate(p):** This primitive allows clusters to be activated independently with probability p by having followers pull the outcome of a p -biased coin flipped by their leader. Each cluster stays activated until the next call of ClusterActivate or until it is explicitly deactivated.
- **ClusterSize:** Each cluster can determine its size in two rounds by each cluster follower sending its ID to the cluster leader, the cluster leader counting the received IDs and all followers pulling this count from the leader.
- **ClusterDissolve(s):** This primitive ensures that all clusters are at least of the given size s and runs in two rounds: In the first round each cluster follower pushes their ID to the leader. The cluster leader uses these messages to determine the cluster size s' . In the second round followers pull a value indicating their new clustering status from their current leader and set their *follow* variable to this value. If $s' \geq s$ the cluster leader responds with its own ID and otherwise responds with ∞ to dissolve the cluster. It also sets its own *follow* variable the same way.
- **ClusterResize(s):** This primitive ensures that all clusters are of size at most $2s$. It runs in two rounds. In the first round each cluster follower pushes their ID to the leader. This allows the leader of any cluster to determine the size s' of its cluster. In the second round followers pull a message from their current leader and update their *follow* variable as follows: The cluster leader (in his mind) re-clusters its followers into $\lfloor \frac{s'}{s} \rfloor$ clusters of equal size (up to one) declaring for each new cluster the node with the largest ID its leader. To announce² and establish this clustering the leader responds to each node's pull request with the IDs of the $\lfloor \frac{s'}{s} \rfloor$ new leaders. Each follower then sets its *follow* variable to the largest leader ID that is at least as large as its own ID. It is easy to see that after a cluster resizing step all clusters have size at most $2s - 1$.
- **ClusterPUSH(msg)/ClusterPULL(msg):** This is a similar primitive to a random PUSH or PULL of a node except on a cluster level. Each cluster leader decides on whether a push or pull is to be performed and what message msg to send or with what message msg to respond if pulled. This decision is then shared within each cluster by followers pulling it from the leader. All nodes in clusters instructed to do now contact a random node and either relay the message of the cluster

²The messages used in this primitive are of size $O(\frac{s'}{s} \log n)$ bits. This calls for ensuring that $\frac{s'}{s} = \Theta(1)$ if a small bit complexity is desired. We note that this complication arises only for address-oblivious algorithms as a non address-oblivious algorithm could simply respond to each follower separately with the ID of its new leader.

leader via a PUSH or request a message via a PULL. All messages received by a node through either a PULL or PUSH get then relayed to their cluster leader (if there is one).

- **ClusterMerge(new leader ID):** This primitive allows to merge clusters in one round. For this all followers simply request the ID of the new leader from their current leader to update their *follow* variable. In addition to answering these pulls the leader also sets its own *follow* variable to the new leader.
- **ClusterShare(.):** This primitive allows clusters to share information within a cluster in two rounds. In the first round, followers who have relevant information send this information to the leader, who aggregates the information. In a second round, followers pull the aggregated information from the leader.

4. GOSSIP IN $O(\log \log N)$ ROUNDS

In this section we present the simplest version of our algorithm to demonstrate the principal ideas needed achieve the optimal $O(\log \log n)$ round-complexity. The resulting algorithm CLUSTER1 is, for sake of simplicity, not tuned to have a good message- or bit-complexity. In particular, a constant fraction of the nodes is sending during most of the $\Theta(\log \log n)$ rounds and the message sizes occurring in the first ClusterResize call is hard to analyze.

4.1 Overview

The idea of CLUSTER1 is simple, and a general outline can be inferred from the procedure calls and comments at the top of Algorithm 1: To obtain an initial clustering, Procedure GROWINITIALCLUSTERS samples a small logarithmic fraction of the nodes to active singleton-clusters. Then a standard PUSH gossip is performed for $\Theta(\log \log n)$ rounds in which these clusters recruit unclustered nodes. The usual exponential growth behavior of the PUSH gossip protocol leads to clusters of size at least $s = C' \log n$ while also ensuring that all but a small constant fraction of the nodes are recruited into some cluster. In the main phase of the algorithm, namely Procedure SQUARECLUSTERS, we repeatedly square this cluster size by merging clusters in the following manner: We start by using ClusterDissolve(s) and ClusterResize(s) to ensure that clusters have sizes between s and $2s$. Next, we determine the new cluster leaders by activating each cluster with probability $1/s$. The inactive clusters are then prompted to join one of the new active clusters via two ClusterPUSHes. This results in each new cluster recruiting $\Theta(s)$ inactive clusters which leads to the desired new cluster size of $\Theta(s^2)$. In particular, the first such PUSH guarantees that each new cluster grows by a factor of $\Theta(s)$, while the second iteration ensures that all old inactive clusters are merged into some new, larger active cluster. The doubly exponential growth of squaring the cluster size with every iteration guarantees that $\Theta(\log \log n)$ iterations are sufficient to obtain clusters of size $\frac{\sqrt{n}}{\log n}$. Procedure MERGEALLCLUSTERS uses two ClusterPUSHes to merge all clusters into the cluster with the smallest ID. With one cluster comprising most of the nodes in the network, Procedure UNCLUSTEREDNODESPULL uses a standard PULL gossip to have the remaining unclustered nodes join the cluster which is known to require only $\Theta(\log \log n)$ rounds. Lastly, with

all nodes in one cluster spreading any desired message can be done via one ClusterShare.

A breakdown of the various steps is given in the detailed procedures of Algorithm 1, using the cluster communication primitives from Section 3.2.

Algorithm 1 CLUSTER1

```

1: GROWINITIALCLUSTERS           ▷ cluster most nodes
2: SQUARECLUSTERS               ▷ repeatedly merge clusters
3: MERGEALLCLUSTERS             ▷ merge clusters to one
4: UNCLUSTEREDNODESPULL        ▷ unclustered nodes join
5: ClusterShare(message)

6: procedure GROWINITIALCLUSTERS
7:   follow ← ∞
8:   With probability  $1/C \log n$  set follow ← ownID
9:   for  $\Theta(\log \log n)$  times do
10:    ClusterPUSH(clusterID)
11:    unclustered nodes: follow ← received ID (if any)

12: procedure SQUARECLUSTERS
13:    $s = C' \log n$ 
14:   ClusterDissolve( $s$ )
15:   repeat
16:     ClusterResize( $s$ )
17:     ClusterActivate( $1/s$ )
18:     for two repetitions do
19:       active: ClusterPUSH(follow)
20:       inactive: ClusterMerge(smallest received ID)
21:        $s \leftarrow \Theta(s^2)$ 
22:   until  $s > \frac{\sqrt{n}}{\log n}$ 

23: procedure MERGEALLCLUSTERS
24:   for two repetitions do
25:     ClusterPUSH(follow)
26:     ClusterMerge(smallest received ID)

27: procedure UNCLUSTEREDNODESPULL
28:   for  $\Theta(\log \log n)$  times do
29:     unclustered nodes PULL follow value

```

4.2 Correctness Proof

In this section we prove the correctness and efficiency of CLUSTER1, which is described in Section 4.1 and Algorithm 1. We break the proof down to several lemmas, which then lead to the main result.

LEMMA 5. *At the end of Procedure GROWINITIALCLUSTERS in CLUSTER1, at least $0.9n$ nodes are contained in clusters of size $C' \log n$ with high probability.*

PROOF. The clustering at the end of the procedure is created by $\Theta(\log \log n)$ rounds of random PUSH gossip starting from cluster centers that were obtained by sampling nodes independently with probability $p = \frac{1}{C \log n}$. By a Chernoff

bound, we initially have the expected $\Theta\left(\frac{n}{C \log n}\right)$ cluster centers occur with high probability. By standard gossip analysis, the set containing all clustered nodes grows by a constant factor in each round, so long as some constant fraction of the system remains unclustered. Furthermore, by choosing $C' \ll C$ at most an arbitrary small constant fraction of nodes account for clusters that are smaller than $C' \log n$, namely about $\frac{n}{C \log n} \cdot C' \log n$ many. In total, one can guarantee that at least a 90% fraction of nodes are good, that is, contained in a large cluster. \square

LEMMA 6. *Every iteration inside Procedure SQUARECLUSTERS in CLUSTER1 squares the cluster size s while ensuring that all clustered nodes are contained in clusters of the new size, with high probability.*

PROOF. Initially, due to Theorem 5 which is followed by ClusterDissolve(s), we can assume that at least $0.9n$ nodes are clustered and divided into clusters of size between s and two $2s$ after the ClusterResize(s) in Line 16 is performed. The same holds from there on by induction. These $\Theta(n/s)$ clusters then get independently activated with probability $1/s$. Since $s \leq \frac{\sqrt{n}}{\log n}$ the expected number of activated clusters is $\Theta(n/s^2) = \omega(\log n)$ and a Chernoff bound shows that indeed $\Theta(n/s^2)$ clusters are activated with high probability. Now, in the first ClusterPUSH/ClusterMerge iteration of Procedure SQUARECLUSTERS, these active clusters PUSH $\Theta(s)$ messages each for a total of $\Theta(n/s)$ messages in order to recruit the $\Theta(n/s)$ inactive clusters. By standard balls and bins analysis, a linear fraction of every set of $s = \Omega(\log n)$ PUSHes of each cluster lands in a unique cluster and result in a merge. Therefore, after the first iteration all active clusters will recruit $\Theta(s)$ inactive clusters and thus grow to size $\Theta(s^2)$. However, there is also a constant fraction of inactive clusters that did not get recruited. The second ClusterPUSH/ClusterMerge iteration ensures that these clusters merge with some active cluster to ensure that all clustered nodes remain clustered. This is indeed the case since there are $\Theta(n)$ total number of nodes contained in the active clusters which PUSH out in the second iteration. Every inactive cluster of size $s = \Omega(\log n)$ is therefore reached and merged into an active cluster with high probability. \square

LEMMA 7. *After Procedure MERGEALLCLUSTERS in Algorithm CLUSTER1, with high probability, there is a single cluster which contains at least $0.9n$ nodes.*

PROOF. Theorem 6 shows that before Procedure MERGEALLCLUSTERS is invoked at least $0.9n$ nodes are contained in clusters of size $\Theta(s)$ with $s \geq \frac{\sqrt{n}}{\log n}$. We now consider the cluster with the smallest ID. The first ClusterPUSH/ClusterMerge iteration of Procedure MERGEALLCLUSTERS will, with high probability, grow this cluster to size $\Omega\left(\frac{n}{\log^2 n}\right)$. The second iteration then recruits all remaining clusters, with high probability. \square

Lastly, it is well known (see, e.g., [10]) that the standard PULL gossip takes only $O(\log \log n)$ rounds to inform the remaining constant or even $1 - 1/\log^{\Theta(1)} n$ fraction of nodes. This is exactly what is used in Procedure UNCLUSTEREDNODESPULL to have all unclustered nodes join the single cluster in the end:

LEMMA 8. *The PULL gossip given in Algorithm CLUSTER1 as Procedure UNCLUSTEREDNODESPULL results in all*

nodes joining the single cluster, with high probability, in $O(\log \log n)$ rounds.

PROOF. It is well known that the standard uniform PULL gossip terminates with high probability in $\Theta(\log \log n)$ if one starts with at least $n/\log^{\Theta(1)} n$ nodes being informed. For completeness we include the analysis here: We look at the fraction x of unclustered nodes and note that in one round any unclustered node (pulling a random node) remains unclustered with probability x . As long as the number of unclustered nodes is still at least $\Theta(\log n)$ a Chernoff bound shows that the quantity x shrinks to at most $2x^2$ with high probability after each pull. This can happen at most for $2 \log \log n$ rounds. Lastly, with high probability, it takes at most a constant number of rounds until each remaining node pulls from a clustered node. This shows that the $\Theta(\log \log n)$ rounds of PULL gossip indeed, with high probability, lead to one cluster containing all nodes. \square

The correctness and efficiency of CLUSTER1 now follows immediately:

THEOREM 9. *Algorithm CLUSTER1 runs for $O(\log \log n)$ rounds and, with high probability, spreads any message to all nodes.*

PROOF. The ClusterShare(message) performed at the end of CLUSTER1 correctly spreads any message to all nodes since according to Theorem 8 all nodes are contained in one single cluster. The running time of CLUSTER1 is furthermore easily seen to be $O(\log \log n)$ rounds. In particular, the loops inside Procedures GROWINITIALCLUSTERS and UNCLUSTEREDNODESPULL run explicitly for $O(\log \log n)$ iterations, while the loop inside Procedure SQUARECLUSTERS squares the parameter s in each iteration, and therefore runs for at most $\Theta(\log \log n)$ iterations as well. To complete the proof, we refer to Section 3.2 for the fact that all cluster primitives performed in these loops can be performed in a constant number of rounds. \square

5. LINEAR MESSAGE COMPLEXITY

In this section we prove Theorem 2 by explaining how the CLUSTER1 algorithm can be tweaked to also have optimal message- and bit-complexity without compromising its optimal $O(\log \log n)$ round-complexity. We call the resulting algorithm CLUSTER2.

5.1 Overview

First, we give the intuition behind achieving a linear message complexity and describe in which aspects CLUSTER2 is different from CLUSTER1. The high-level recipe of CLUSTER2 is the same as CLUSTER1. The key idea for achieving the optimal message-complexity is in modifying Procedures GROWINITIALCLUSTERS and SQUARECLUSTERS to work with only $\Theta(n/\log n)$ clustered nodes. This guarantees that even if all clustered nodes actively transmit in these rounds, only $o(n)$ messages are used for these parts. After merging all clusters, we are left with one large cluster of size $\Theta(n/\log n)$. While the standard PULL gossip would still finish in $O(\log \log n)$ rounds this would again take to many messages (of unclustered nodes PULLing unclustered nodes). We therefore PUSH the cluster ID to $\Theta(n)$ nodes first before doing the final PULL phase. With so many informed nodes each node expects to PULL at most a constant

number of times and it is easy to see that in total at most $O(n)$ messages are used with high probability.

All in all a much more delicate control on the number of clusters, their size and the number of clustered nodes is required. This is also important to provably guarantee an optimal bit-complexity. The subtle problem here is that a ClusterResize(s) on a cluster of size $\omega(s)$ leads to messages of size $\omega(\log n)$ (see footnote in Section 3.2). This control is achieved as follows:

To control the number of nodes clustered in the Procedures GROWINITIALCLUSTERS we have clusters measure their growth. More precisely before and after each PUSH in the initial recruiting phase each cluster checks its size. If a cluster has size at least $\Omega(\log^3 n)$ but grew by less than a factor of $2 - \frac{1}{\log n}$ then it stops recruiting. A Chernoff Bound shows that no cluster will stop recruiting if less than $n/2 \log n$ nodes are clustered. On the other hand, similar arguments as for CLUSTER1 show that indeed $\Theta(n/\log n)$ nodes will get recruited of which most lie in clusters of size $s = \Omega(\log^3 n)$. A ClusterResize($C' \log n$) called for every large cluster in every iteration furthermore ensures that no individual cluster gets too big. For the Procedure SQUARECLUSTERS we then again merge clusters. Due to the smaller number of clustered nodes most cluster requests now go to unclustered nodes and get ignored. Still, any iteration starting with clusters of size s ends up with larger clusters of size $\Theta(s^2/\log n) = \omega(s^{1.5})$. Growing the cluster sizes from $s = \Theta(\log^3 n)$ to $\Omega(n/\log^2 n)$ therefore still takes only $O(\log \log n)$ rounds. Chernoff bounds over the larger $\Omega(\log^3 n)$ sized clusters still guarantee that the new cluster size can be predicted up to a constant with high probability which ensures that no ClusterResize(s) is applied to a too large (or too small) cluster.

The formal description of CLUSTER2 is given as Algorithm 2.

5.2 Correctness Proof

In this section we prove Theorem 2 by showing that CLUSTER2 spreads a message with the desired optimal round-, message- and bit-complexity.

The proof follows along the same line as Theorem 9 while including the extra arguments given in Section 5.1. We again break the proof into several lemmas, which then lead to the main result.

LEMMA 10. *If at most $n/f(n)$ nodes are clustered, then a ClusterPUSH by some cluster X , where $|X| \geq \Omega(\log n)$, grows X by factor at least $(2 - \Theta(1/f(n)))$ with high probability.*

PROOF. Each PUSH request of one of the $|X|$ nodes is hits a unique unclustered node with probability of approximately $(1 - \frac{1}{f(n)})$. The expected number of nodes recruited for cluster X is therefore $(1 - \frac{1}{2f(n)})|X|$ for a total number of $|X| + (1 - \frac{1}{2f(n)})|X|$ and a growth factor of $(2 - \Theta(1/2f(n)))$. Since the recruitment success of different nodes are essentially independent of each other applying a Chernoff type bound shows that the growth factor is close to the expected and in particular is $(2 - \Theta(1/f(n)))$, with high probability. \square

LEMMA 11. *Procedure GROWINITIALCLUSTERS of CLUSTER2 leads to $\Theta(n/\log n)$ clustered nodes that are contained*

Algorithm 2 CLUSTER2

```
1: GROWINITIALCLUSTERS
2: SQUARECLUSTERS
3: MERGEALLCLUSTERS           ▷ as in Algorithm 1
4: BOUNDEDCLUSTERPUSH
5: UNCLUSTEREDNODESPULL       ▷ as in Algorithm 1
6: ClusterShare(message)

7: procedure GROWINITIALCLUSTERS
8:   follow  $\leftarrow \infty$ 
9:   With probability  $1/C \log^4 n$  set follow  $\leftarrow$  ownID
10:  ClusterActivate(1)
11:  for  $\Theta(\log \log n)$  times do
12:    active clusters: ClusterPUSH(follow)
13:    unclustered nodes: follow  $\leftarrow$  any received ID
14:    if ClusterSize  $\geq C' \log^3 n$  then
15:      if cluster growth factor  $< (2 - 1/\log n)$  then
16:        deactivate cluster
17:      else
18:        ClusterResize( $C' \log^3 n$ )

19: procedure SQUARECLUSTERS
20:    $s = C' \log^3 n$ 
21:   ClusterDissolve( $s$ )
22:   repeat
23:     ClusterResize( $s$ )
24:     ClusterActivate( $1/s$ )
25:     for two repetitions do
26:       active: ClusterPUSH(follow)
27:       inactive: ClusterMerge(random received ID)
28:      $s \leftarrow \Theta(s^2/\log n)$ 
29:   until  $s > \frac{\sqrt{n}}{\log^2 n}$ 

30: procedure BOUNDEDCLUSTERPUSH
31:   ClusterActivate(1)
32:   for  $\Theta(\log \log n)$  rounds do
33:     ClusterPUSH(follow)
34:     unclustered nodes: follow  $\leftarrow$  any received ID
35:     ClusterSize
36:     if cluster growth factor  $< 1.1$  then
37:       deactivate cluster
```

in clusters of size $(1 + \Theta(1))(C' \log n)$, with high probability, while at most $O(n)$ messages of $\Theta(\log n)$ bits each are sent.

PROOF. For the initial sampling of singleton clusters, for which the sampling probability is $1/C \log^4 n$, a Chernoff bound shows that $\Theta(\frac{n}{C \log^4 n})$ such singleton clusters are created, with high probability. We would like to argue as in Theorem 5 that the growth of the set of clustered nodes is exponential until at least $\Omega(n/\log n)$ nodes are clustered. We need to show that the newly introduced size-control does not prevent this exponential growth. By Theorem 10, as long as at least $n(1 - O(1/\log n))$ nodes are unclustered each cluster will grow by a factor of at least $2 - 1/\log n$, with high probability. Hence, the size control mechanism does not inhibit the exponential growth until at least $\Theta(n/\log n)$ nodes are clustered. Since at most about $\frac{n}{C \log^4 n} \cdot (C' \log^3 n)$ nodes account for clusters that are smaller than $C' \log^3 n$, indeed at least $\Theta(n/\log n)$ nodes are contained in large clusters.

To guarantee that the message complexity incurred by the initialization is small we note that only clustered nodes are transmitting. This leads to at most $O(n \log \log n / \log n) = O(n)$ messages sent. Since every cluster can grow by at most a factor of two in each iteration it is furthermore clear that all message sizes used in a ClusterResize contain at most two IDs. This results in messages of size $\Theta(\log n)$, as claimed. \square

LEMMA 12. *Before and after every iteration inside Procedure SQUARECLUSTERS in CLUSTER2, $\Theta(n/\log n)$ nodes are contained in clusters of size at least s , with high probability. Furthermore, the total number of messages sent during Procedure SQUARECLUSTERS is $O(n)$, each containing $\Theta(\log n)$ bits.*

PROOF. We prove the lemma by induction. The base case is already given by Theorem 11. Here, we show that each iteration of the loop increases the cluster size from s to $s' = \Theta(\frac{s^2}{\log n})$ while keeping all clustered nodes clustered. In the first iteration each active cluster PUSHes to $\Theta(s)$ random nodes and hits with high probability $\Theta(s/\log n)$ inactive clusters (alone). Once these clusters merge, any active cluster therefore grows to $s' = \Theta(\frac{s^2}{\log n})$ as desired. The second PUSH, performed by the larger clusters which cover $\Theta(n/\log n)$ nodes, guarantees that no inactive cluster of size s is left behind without merging to a new larger cluster of size at least s' . Furthermore, since inactive clusters randomly join the large active clusters during this second iteration no active cluster grows by more than a constant factor. The new clusters are therefore all of the desired size $s' = \Theta(\frac{s^2}{\log n})$ while all clustered nodes remain clustered. Lastly, since $s > \log^3 n$ we have that $s' = \omega(s^{1.5})$ which shows that the number of iterations until the final cluster size of $\Omega(\frac{\sqrt{n}}{\log^2 n})$ is reached is at most $\Theta(\log \log n)$.

All in all, only $\Theta(n/\log n)$ nodes are communicating in each of the $O(\log \log n)$ rounds, which leads to the desired $O(n)$ messages-complexity. Regarding the bit-size of these messages, the only thing to worry about is the calls to ClusterResize(s). However, since the cluster sizes are, up to constants, exactly what they are resized to, ClusterResize(s) produces only $\Theta(\log n)$ bit messages containing $\Theta(1)$ IDs. \square

Theorem 7 applies without change to CLUSTER2 and we do not repeat it.

LEMMA 13. *Procedure BOUNDEDCLUSTERPUSH expands the single large cluster to $\Theta(n)$ nodes in $O(\log \log n)$ rounds, with high probability, while at most $O(n)$ messages are sent.*

PROOF. At the beginning of the procedure, we have from Theorem 7 that $\Theta(n/\log n)$ belong to a single giant cluster. From here, we use a standard analysis of PUSH-PULL gossip: So long as at least $.1n$ nodes are unclustered, during each iteration the number of clustered nodes grows exponentially (and therefore for at most $O(\log \log n)$ rounds) until a large constant fraction of the nodes is clustered. By Theorem 10, the cluster detects this event successfully. Since the number of messages sent out in every round is exponentially increasing until it stops the total number of messages form a geometric sum which sums to $O(n)$ messages sent in total. \square

Putting these claims together, it follows that CLUSTER2 has a bit-complexity of $O(n \log n + nb) = O(nb)$ and message complexity of $O(1)$ messages per node on average while maintaining the optimal $O(\log \log n)$ round complexity of CLUSTER1. This proves Theorem 2.

6. THE $\Omega(\log \log N)$ ROUND LOWER BOUND

This section is dedicated to prove the $\Omega(\log \log n)$ round-complexity lower bound of Theorem 3.

We first remark that the simple lower bound provided in Section 7 as Theorem 16 does not give any super constant number of rounds for the most interesting setting of $\Delta = n^{\Theta(1)}$. Here we show that the round-complexity of broadcasting a message has to be at least $0.99 \log \log n$ with high probability. This holds for any Δ and remains true even if one ignores many restrictions of the model regarding message sizes, address-obliviousness, or the number of known nodes that can be contacted per round. The main implication of this lower bound is that both CLUSTER1 and CLUSTER2 achieve the optimal asymptotic running time.

We need some notation: Instead of each node sampling (up to) one random contact in each round we assume that these nodes are sampled in advance. In particular, let $u_{v,t}$ be the uniformly random node given to node v if it samples a node at time t . Based on these samples we define G_t to be the union of all edges (node-pairs) that are potentially sampled at time t , that is, $G_t = (V, E_t)$ with $E_t = \{\{u, v\} \mid \exists v : u = u_{v,t}\}$.

LEMMA 14. *For every $t = 0, 1, \dots$ let K_t be the graph indicating which node knows whom at the beginning of round t . It holds that: $K_0 = \emptyset$; $K_{t+1} \subseteq (K_t \cup G_{t+1})^2$, and therefore also that:*

$$K_t \subseteq \left(\bigcup_{i=1}^t G_i \right)^{2^t}.$$

Here G^j is the graph that connects a node u to any node for which there is a path of length at most j in G .

PROOF. The first equation is simply stating that nodes do not know each other initially. To see the second equation we note that there are two ways nodes can learn about new nodes. The first one is by contacting a random node. This only adds edges in G_{t+1} to the knowledge graph K_t . By initiating a contact along an edge in G_{t+1} or along an edge in K_t corresponding to a node known from previous rounds, nodes can also indirectly learn about nodes known to nodes they know. Even if all such “neighbors” are contacted and every node responds with everything they know (which would result in many message of very large size), a

nodes does not learn about more nodes than its 2-hop neighborhood. This is exactly what is stated in the second equation. The last equation now follows from the other two by induction and taking into account that for any two graphs G, G' and any $j, j' > 1$ we have $G^j \cup G'^j \subseteq (G \cup G')^j$ and of course $(G^j)^{j'} = G^{jj'}$. \square

We can now prove Theorem 3:

THEOREM 15 (THEOREM 3 RESTATED). *Suppose we allow unlimited message sizes, non address-oblivious communication, and have no restriction on how many communications a node can be involved in or even to how many known nodes a node contact in one round. Still, any algorithm running in less than $T = \log \log n - \log \log \log n - \omega(1)$ rounds fails to spread a message that starts at one node u to all nodes with high probability.*

PROOF. First we note that spreading a message starting at one node u can be used to elect u as a cluster leader by simply attaching its ID to the message spread. This could then be used to make K_T the complete graph by informing all nodes about each other in at most two additional rounds, at least if the message size is unbounded. We show that this is, with high probability, impossible in less than T rounds by using Theorem 14. In particular, it certainly holds that $K_T \subseteq K'^{2^T}$ where $K' = (\bigcup_{i=1}^T G_i)$. Note that K' is simply a random graph that is created by each node sampling T neighbors independently and adding an undirected edge to those neighbors. It is clear that K_T is a complete graph if and only if the diameter of K' is at most 2^T . It is easy to see that for $T = \log \log n - \log \log \log n - \omega(1)$ the probability that such a random graph has diameter $o(\log n / \log \log n)$ is extremely small which completes the proof. An easy way to get such a high probability result is to say that the probability that there is a node in K' with degree $\Delta = \omega(\log n)$ is at most $n^{-\omega(1)}$ and any graph with maximum degree $\Theta(\log n)$ necessarily has a diameter of $\Omega(\log n / \log \log n)$. \square

7. BOUNDED COMMUNICATIONS

In this section we discuss and prove tradeoffs between the running time of any gossip algorithm and the maximum number of communications a node can be involved in per round.

In the standard gossip model which is also employed in this paper, each node can only initiate one PUSH or PULL communication per round. Still, a node might be involved in multiple communications initiated by other nodes. The possibility of this happening is present in the models of *all* prior gossip works (discussed in Section 1). In few cases, in particular the random phone call model with an underlying complete graph, having a large number of communications per round in one node is not crucial. For example, the running time of the uniformly random gossip on the complete graph remains $O(\log n)$ even if each node only accepts one communication per round (but the result of [12] which concentrates on the leading constants and lower order terms change). For most gossip works however such a concentration of communications is unavoidable. For example the works [4–6] which generalize the random phone call model to arbitrary topologies show that gossip is fast on any topology with good vertex or edge expansion. In particular, they imply that gossip completes in $\log^{O(1)} n$ rounds on a star

graph because the star graph has a constant edge expansion $\Phi = O(1)$ and constant vertex expansion $\alpha = O(1)$. This is obviously not possible without the center node talking to at least $\Delta = n/\log^{O(1)} n$ nodes per round on average. Similarly, all results that use direct addressing crucially rely on each node being able to respond to many nodes and many of them can have one node communicate with (almost) all other nodes in the network in one round. This is also the case for CLUSTER1 and CLUSTER2. In fact our entire set of cluster-operations are based on having a cluster-center coordinating $\omega(1)$ cluster-followers in a single communication step.

While in some settings having $\Delta = n$ may be acceptable, in other settings one would like to have Δ to be not much more than a small polynomial or even (significantly) less. In what follows we explicitly address this parameter and design optimal algorithms for any Δ .

We first remark that having a super-constant Δ is unavoidable if one wants to achieve a sub-logarithmic round complexity. In particular, it is easy to see that any global communication with the restriction that no node responds to more than Δ requests per round takes at least $\log n/\log \Delta$ rounds:

LEMMA 16. *Any algorithm in which no node participates in more than Δ communications per round requires at least $\log n/\log \Delta$ rounds to inform all nodes of a message.*

PROOF. Since each informed node can participate in at most Δ communications the number of informed nodes cannot grow by more than a Δ factor per round. Inductively the the number of informed nodes in round r is therefore at most Δ^{r-1} and $\log n/\Delta$ rounds are needed to bring this quantity to n . \square

In the remainder of this section we will show that the algorithmic ideas developed in this paper can be used to achieve any point on the this tradeoff curve between Δ and the required number of rounds while also maintaining the optimal message- and bit-complexity. We do this in two steps. We first define a Δ -clustering and show that any such clustering easily allows Δ -bounded communications in the desired $\log n/\log \Delta$ round complexity by running simple and robust gossip algorithms on top of this clustering. We then show that a Δ -clustering can be computed in $\Theta(\log \log n)$ rounds:

DEFINITION 1. *A Δ -clustering is a clustering in which every node is clustered and all clusters have a size between Δ and 2Δ .*

We will show in Theorem 17 that such a Δ -clustering indeed allows to efficiently spread information in the desired $O(\log n/\log \Delta)$ rounds using the CLUSTERPUSH-PULL protocol given in Algorithm 3. For sake of simplicity we focus on the case $\Delta = \log^{\omega(1)} n$ since for smaller Δ the improvement of $\log \Delta$ is at most $O(\log \log n)$ and therefore essentially negligible. Similarly, in this section we assume for convenience that $\Delta < n^{0.9}$ as otherwise one can might as well set $\Delta = n$ and run CLUSTER2.

LEMMA 17. *For any bound $\Delta = \log^{\omega(1)} n$ the algorithm CLUSTERPUSH-PULL(Δ) given in Algorithm 3 spreads a message to all nodes, with high probability. It furthermore uses only $O(\frac{\log n}{\log \Delta})$ rounds and $O(n)$ messages once the Δ -clustering is computed.*

Algorithm 3 CLUSTERPUSH-PULL(Δ)

- 1: Compute Δ -Clustering
 - 2: ClusterShare(message)

 - 3: **for** $\Theta(\frac{\log n}{\log \Delta})$ iterations **do**
 - 4: **if** newly informed clusters **then**
 - 5: ClusterPUSH(message), ClusterShare(message)

 - 6: uninformed nodes: PULL message from a random node

 - 7: ClusterShare(message)
-

PROOF. We first address the round and message complexity and then prove correctness:

The claimed round complexity of $O(\frac{\log n}{\log \Delta})$ is easily verified since the only loop of the algorithm in Line 3 takes 3 rounds per iteration and runs for $\Theta(\frac{\log n}{\log \Delta})$ iterations. All other instructions take only $O(1)$ rounds. It is also easy to see that each node processes at most $O(1)$ messages over the course of the algorithm.

Next we prove correctness, that is, that CLUSTERPUSH-PULL spreads any message to all nodes with high probability. The first ClusterShare (Line 5) spreads the message to $\Delta = \omega(\log n)$ nodes. From there on, so long as at most $n/\Omega(\Delta)$ nodes are informed, each ClusterPUSH and ClusterShare in the main loop (of Line 3) increases the number of informed nodes with high probability by a $\Delta/O(1)$ factor. Hence, the set of informed nodes grows exponentially until it reaches $n/\Delta > (\Delta/O(1))^{\Theta(\frac{\log n}{\log \Delta})}$ nodes.

To see that this suffices for the ClusterPULL/ClusterShare in Lines 6-7 to inform all remaining nodes with high probability we focus on one uninformed node. This node gets informed during the ClusterPULL if at least one of the nodes in its cluster PULLs the message from a node in Line 6. Since there are at least $\Delta = \log^{\omega(1)} n$ such nodes of which each has an independent chance of $1/\Delta$ to PULL from an informed node this happens with high probability.

This shows that CLUSTERPUSH-PULL indeed informs all nodes about the message with high probability using only $O(\frac{\log n}{\log \Delta})$ rounds and $O(n)$ messages. \square

What remains to show is that a $\Theta(\Delta)$ -clustering can be computed efficiently. The process for generating the clustering is essentially a generalization of the CLUSTER2 algorithm, where we stop growing clusters at around size Δ . Then for additional $O(\log \log n)$ rounds, we join all unclustered nodes to existing clusters. All the while, we never connect any leader to more than Δ followers, nor transmit more than $O(n)$ messages. The detailed pseudo-code description of the resulting algorithm CLUSTER3(Δ) is given in Algorithm 4.

THEOREM 18. *For any given $\Delta = \log^{\omega(1)} n$ the algorithm CLUSTER3(Δ), with high probability, computes a $\Theta(\Delta)$ -clustering using $O(\log \log n)$ -rounds and $O(n)$ messages while no node communicates with more than Δ nodes in any round.*

PROOF. As before, the Procedures GROWINITIALCLUSTERS and SQUARECLUSTERS result in $\Theta(n/\log n)$ clustered nodes grouped into clusters of size $\Theta(s)$. Here the cluster

Algorithm 4 CLUSTER3(Δ) with $\Delta = \log^{\omega(1)} n$

```
1: GROWINITIALCLUSTERS           ▷ as in Algorithm 2
2: SQUARECLUSTERS(to size  $s \geq \sqrt{\Delta \log n / C''}$ )
                                   ▷ as in Algorithm 2

3: MERGECLUSTERS
4: BOUNDEDCLUSTERPUSH
5: UNCLUSTEREDNODESPULL         ▷ as in Algorithm 1
6: ClusterResize( $\Delta / C''$ )

7: procedure MERGECLUSTERS
8:   ClusterActivate( $10 \frac{s}{\Delta / C''}$ )
9:   active: ClusterPUSH(clusterID)
10:  inactive: ClusterMerge(random received ID)

11: procedure BOUNDEDCLUSTERPUSH
12:   ClusterActivate(1)
13:   for  $\Theta(\log \log n)$  rounds do
14:     ClusterResize( $\Delta / C''$ )
15:     ClusterPUSH(follow)
16:     unclustered nodes: follow  $\leftarrow$  any received ID
17:     ClusterSize
18:     if cluster growth factor  $< 1.1$  then
19:       deactivate cluster
```

size is $s = \Omega(\sqrt{\Delta \log n / C''})$. The proof for this follows the same way as for CLUSTER2, and we do not repeat the arguments here.

Procedure MERGECLUSTERS of CLUSTER3 is a bit different from CLUSTER2, though the techniques are similar. It activates every cluster with probability $10 \frac{s}{\Delta / C''}$ followed by a ClusterPUSH of the active clusters. Since Δ is small enough a Chernoff bound shows that roughly a $\frac{10C''}{\Delta}$ fraction of clusters is activated. Since $\Delta = \log^{\omega(1)} n$ each inactive cluster gets furthermore PUSHed by multiple active clusters to merge. Importantly, and differently to other merging phases, inactive clusters choose one to merge with among these clusters uniformly at random. A Chernoff type-bound shows that this way inactive clusters get assigned to active clusters evenly, i.e., each active cluster gets chosen by $\Theta(\frac{\Delta / C''}{s})$ inactive clusters and will therefore grow to size about $\Delta / 10C''$.

Subsequently, these clusters then recruit unclustered nodes in the BOUNDEDCLUSTERPUSH call. After each recruiting step clusters measure whether they grew by at least a factor of 1.1 and resize. As before this guarantees that all but a small constant fraction of the nodes joins a cluster while the number of messages is $O(n)$. Resizing active clusters guarantees that no cluster gets too big and no leader has to communicate too much.

As a last step, we run UNCLUSTEREDNODESPULL on the unclustered nodes and make them join the first cluster they PULL from. Up to constants every unclustered node joins a uniformly random cluster. Even without any size control every cluster grows by at most a small constant in this round. The final ClusterResize(Δ / C'') ensures that all clusters have the same size (up to a factor of two) as desired by a $\Theta(\Delta)$ -clustering. \square

Acknowledgments

We thank Chen Avin and Robert Elsässer for sharing an early version of [1] and for interesting discussions.

8. REFERENCES

- [1] C. Avin and R. Elsässer. Faster Rumor Spreading: Breaking the $\log n$ Barrier. In *Proceedings of the International Symposium on Distributed Computing*, DISC '13, pages 209–223, 2013.
- [2] K. Censor-Hillel, B. Haeupler, J. Kelner, and P. Maymounkov. Global Computation in a Poorly Connected World: Fast Rumor Spreading with no Dependence on Conductance. In *Proceedings of the ACM Symposium on Theory of Computing*, STOC '12, pages 961–970, 2012.
- [3] K. Censor-Hillel and H. Shachnai. Fast information spreading in graphs with large weak conductance. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 440–448, 2011.
- [4] G. Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *Proceedings of the 28th International Symposium on Theoretical Aspects of Computer Science*, STACS, pages 57–68, 2011.
- [5] G. Giakkoupis and T. Sauerwald. Rumor spreading and vertex expansion. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pages 1623–1641, 2012.
- [6] B. Haeupler. Analyzing network coding gossip made easy. In *Proceedings of the ACM Symposium on Theory of Computing*, STOC, pages 293–302, 2011.
- [7] B. Haeupler. Simple, fast and deterministic gossip and rumor spreading. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 705–716, 2013.
- [8] B. Haeupler, G. Pandurangan, D. Peleg, R. Rajaraman, and Z. Sun. Discovery through gossip. In *Proceedings of the 24th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA, pages 140–149, 2012.
- [9] M. Harchol-Balter, T. Leighton, and D. Lewin. Resource discovery in distributed networks. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, PODC'99, pages 229–237, 1999.
- [10] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking. Randomized rumor spreading. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, FOCS, pages 565–574, 2000.
- [11] S. Kutten, D. Peleg, and U. Vishkin. Deterministic resource discovery in distributed networks. *Theory of Computing Systems*, 36(5):479–495, 2003.
- [12] B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, pages 213–223, 1987.