

Confidential Gossip

Chryssis Georgiou
University of Cyprus
chryssis@cs.ucy.ac.cy

Seth Gilbert
National University of Singapore
seth.gilbert@comp.nus.edu.sg

Dariusz R. Kowalski
University of Liverpool
D.Kowalski@liverpool.ac.uk

Abstract—Epidemic gossip has proven a reliable and efficient technique for sharing information in a distributed network. Much of this reliability and efficiency derives from processes collaborating, sharing the work of distributing information. As a result of this collaboration, processes may receive information that was not originally intended for them. For example, some process may act as an intermediary, aggregating and forwarding messages from some set of sources to some set of destinations.

But what if rumors are *confidential*? In that case, only processes that were originally intended to receive the rumor should be allowed to learn the rumor. This blatantly contradicts the basic premise of epidemic gossip, which assumes that processes can collaborate. In fact, if only processes in a rumor’s “destination set” participate in gossiping that rumor, we show that high message complexity is unavoidable. A natural approach is to rely on cryptography, for example, assuming that each process has a well-known public-key that can be used to encrypt the rumor. In a dynamic system, with changing sets of destinations, such a process seems potentially expensive.

In this paper, we propose a scheme in which each rumor is broken into multiple fragments using a very simple coding scheme; any given fragment provides no information about the rumor, while together, the fragments can be reassembled into the original rumor. The processes collaborate in disseminating the rumor fragments in such a way that no process outside of a rumor’s destination set ever receives all the fragments of a rumor, while every process in the destination set eventually learns all the fragments. Notably, our solution operates in an environment where rumors are dynamically and continuously injected into the system and processes are subject to crashes and restarts. In addition, the scheme presented can tolerate a moderate amount of collusion among curious processes without too large an increase in cost.

Keywords—Confidentiality, Collusion, Randomized gossip, Fault-tolerance, Dynamic rumor injection, Message complexity.

I. INTRODUCTION

Collaboration is at the heart of distributed computing: when a network of devices cooperates to solve a problem, the resulting computation is often more robust and more efficient than if each device had worked independently. A classic example of the benefits of collaboration can be found in the paradigm of epidemic gossip. Consider, for example, a set of n devices that want to share information. If each device communicates independently with the other devices in the network, then the message complexity for the protocol may be $O(n^2)$. By contrast, if the devices collaborate to

share the information, each communicating with a small number of random devices in each round, then the message complexity for the protocol can be reduced to $O(n \log n)$. At the same time, the resulting protocol is quite robust, tolerating a constant fraction of the devices crashing.

Yet there are some drawbacks to collaboration. One significant cost is *privacy*: by collaborating with other devices to solve a problem, it is often the case that private information is divulged. Consider again n devices that want to share information—however the information is potentially *confidential* and should only be shared among specified groups of recipients. For example, a user may want to share an engineering blueprint with her colleagues, but not with her competitors. Or a psychiatrist may want to send an e-mail to a group of patients, but not to everyone. Unfortunately, standard distributed protocols for efficiently sharing information do not satisfy these requirements. For example, if the users rely on epidemic gossip to distribute their information, then all confidentiality is lost: every device in the system may learn every piece of information.

In this paper we consider protocols designed to tolerate *honest, but curious* processes. This concept has attracted considerable attention as a model of processes in distributed applications that need limited anonymity and privacy, c.f., [29], [13] (see more in related work below). It is not our protocol’s intent to be secure against truly malicious parties: for data that must be kept secure in all circumstances, a more expensive solution is needed. For everyday transactions, however, where privacy is desired, we can ensure that no process ends up in possession of information that it is not intended to learn. Moreover, we can achieve this additional at relatively limited cost (in terms of message complexity), even if a moderate number of the participants may be colluding (i.e., sharing information).

Results. Thus the question we ask in this paper is whether we can achieve the benefits of collaboration—i.e., robustness and efficiency—without sacrificing confidentiality. We focus our attention on the problem of *Continuous Gossip*, a long-lived version of information sharing (introduced in [12]) that has three notable properties: (1) any process can inject a *rumor* at any time; (2) each rumor specifies a set of recipients that should receive the rumor; and (3) each rumor has a deadline specifying by when it should be received.

In this paper, we present a continuous gossip protocol that guarantees all of the following desirable attributes: **Confidentiality**: Only the specified recipients of a rumor learn the contents of the rumor, even if processes outside the specified recipients may collude. **Timeliness**: Every rumor is delivered by the deadline. **Efficiency**: The maximum *per-round message complexity*, with high probability, and in the absence of collusions, is $O((n^{1+48/\sqrt{dmin}} + n^{1+6/\sqrt[6]{dmin}}) \text{polylog } n)$, where $dmin$ is the shortest deadline of any *active* rumor (that is, a rumor whose deadline has not expired); note that for rumors with deadline of $\Omega(\log^6 n)$, this results in per-round message complexity of $O(n \text{polylog } n)$. When up to τ processes may collude then we show that the maximum per-round message complexity is increased by a factor τ^2 . **Robustness**: Processes may crash and restart at any time; there is no bound on the number of crashed processes at any given time. Moreover, failures are **adaptive**: they may depend on the execution and the random choices made by the individual processes.

Our approach. The major challenge underlying confidential gossip is reconciling the need for collaboration to achieve efficiency, and the inherent loss of confidentiality created by collaboration. At first glance, it seems that only recipients of a rumor can help in its dissemination. Yet, if we limit all information regarding the rumor to its recipients, then it is impossible to achieve good message complexity. As we show in Theorem 1, if we limit messages regarding a rumor ρ to the *destination set* $\rho.D$, then the per-round message complexity is $\Omega(n^{(3/2)-\varepsilon/dmax})$, for any $\varepsilon > 0$ ($dmax$ is the longest deadline of any active rumor). Thus it seems that confidentiality and efficiency are inherently at odds.

We circumvent this seeming impossibility via a simple insight: each rumor can be divided into multiple independent fragments; each fragment provides no information regarding the original rumor, and yet together they can be combined to re-assembled the original rumor. (This is the basic idea underlying *cryptographic secret sharing* [32], [34], though we require only the simplest instantiation of this idea.) All the processes in the system can now collaborate to distribute the rumor fragments, as long we as ensure that no process collects all the fragments. (In fact, we can rely on existing gossip protocols as a black box, as long we restrict their communication to processes that are allowed to receive the specific message fragments.) In this way, we gain the benefits of collaboration without sacrificing confidentiality.

A second challenge we address is the possibility that failures are not independent and history-oblivious. We assume that processes may crash and restart at any time, and we model failures as being caused by an *adaptive* and omniscient adversary that can fail processes based on the random choices made by the protocol. For example, every time a source sends a rumor (or rumor fragment) to another process, the adversary may choose to immediately crash

that recipient, entirely preventing the dissemination of that rumor. We address this challenge by having processes collaborate, exchanging *metadata* that contains no information on rumors. This information allows processes to determine which other processes are failed (or have failed recently), and which processes are being isolated (both in terms of sending and receiving information). Using this metadata, processes can target their messages better, and processes can adjust the number of messages they are sending. By collaborating on metadata, rather than rumors, processes can still overcome an adaptive adversary without giving up confidentiality. (While some information is leaked via the metadata, we discuss in Section VI how to avoid this problem.)

Alternative approaches. There are several possible cryptographic approaches to solving the problem of confidential gossip, many of which exist under the rubric of *multicast security* (e.g., [5], [11], [24], [27], [30], [32]). If a system must be secure against truly malicious parties (i.e., not simply “honest-but-curious” processes), then these cryptographic solutions are the only method of achieving confidentiality.

The basic idea, in many cases, is that each process holds some subset of the cryptographic keys; by encrypting the message with appropriate subsets of the keys, the sender can ensure that the message can only be decrypted by the intended recipients. For example, the processes may be arranged as leaves on a binary tree, where each internal node of the tree contains a cryptographic key; each process is given access to every key found on the root-to-leaf path ending at the leaf owned by the process. When the destination set of a rumor aligns well with the grouping of processes in the tree, such a scheme can be quite efficient; when the destination set contains processes distributed throughout the leafset, then such a scheme can be quite expensive.

Many such solutions (e.g., [2], [26], [33], [35]) focus on a single source communicating confidentially with a single group of processes. The source establishes a shared key with the group, and then updates it as the group changes. Often, a tree-like scheme (as above) is used to make the re-keying more efficient.

In general, the cryptographic solutions will be more efficient when the groupings are stable. That is, when some processes want to communicate with a fixed set of destinations, these cryptographic solutions can be made quite efficient by ensuring that the fixed set of processes share a single cryptographic key. Even when there are occasional changes to the destination set, such solutions work reasonably well. However, we are not aware of any sub-quadratic, in terms of message complexity, cryptographic approach to guarantee confidential gossip when the groups are changing rapidly, or when there are no fixed groups, i.e., when each rumor has a different destination set. In many cases, the best solution appears to be encrypting the message individually

for each process in the destination set, thereby significantly increasing the amount of data to be sent. Furthermore, there is the question on how efficient secret key maintenance would be in the presence of dynamic crashes and restarts, especially when restarted processes have no memory of the computation prior to restarting (as assumed in our model). As we show, our confidential gossip protocol is efficient even under such dynamic adverse conditions.

Other related work. The gossip problem has frequently been considered in relation to random, epidemic communication (e.g., [10], [17], [18], [19]). In this context, the problem is also known as *rumor spreading* and the protocols usually use a simple epidemic approach: each process periodically sends its rumor—along with any new rumors it has learned—to another randomly selected process. This approach can lead to efficient rumor dissemination while tolerating benign failures ([17]).

The gossip problem has been considered in a variety of fault-prone environments, ranging from crash failures to malicious/Byzantine ones (e.g., [6], [15], [21], [23], [25]). The survey by Pelc [31] together with the book by Hromkovic et al. [16] overview solutions for the gossip problem in fault-prone distributed networks.

Another line of work related to ours is the one considering the problem of constructing scalable overlays of topic-based Publication/Subscribe systems (e.g., [1], [7], [8], [28]). The aim is to design an overlay network for each pub/sub topic, so that for each topic, the subgraph induced by the nodes interested in the topic will be connected; such overlays are called *topic-connected*. New events for each topic can then be routed from publishers only to interested subscribers using such topic-connected overlays. If destination sets are viewed as topics, then a topic-connected overlay could provide a confidential way of distributing a rumor to its destination set. Unfortunately, we don't know how to maintain topic-connected overlays in a dynamic setting (i.e., for changing destination sets), and even the static case is NP-complete [8], [28]. Theorem 1 (Section III) effectively implies that topic-connected overlays cannot be used to support efficient confidential gossip.

The *honest-but-curious* model, also referred as the *semi-honest* model [4] is a standard cryptographic adversarial model [13]. This model has been widely considered in the problem of multi-party privacy-preserving computation of some function [36], [13]. The usual demand is for the function to be computed collectively by the computing entities without leaking any information about the entities' inputs, except that revealed by the algorithm's output. Various computations have been considered in the context of computations on sets, such as set union, intersection, element reduction (c.f., [20]), on graphs (c.f., [4]), in the area of data mining (c.f., [22], [3]), majority voting (c.f., [29]), as well as private predicate computation in mobile population

protocols [9].

Our solution to the confidential gossip problem can be viewed as a tool in the process of computing these functions when the privacy of inputs (in the form of rumors) could be kept within groups of processes (i.e., certain rumors would have as their destination set a specific group). For example, a number of group of social networking websites, wishing to efficiently calculate aggregate statistics such as degrees of separation and average number of acquaintances without compromising the in-group privacy, could use as a building block our confidential gossip algorithm.

Paper organization. In Section II we present the model of computation and the confidential continuous gossip problem. In Section III we show that if only the processes of a rumor's destination set collaborate in disseminating the rumor, then high message complexity is unavoidable. In Section IV we present and analyze an efficient randomized algorithm for confidential continuous gossip assuming no collusion. In Section V we show the effect of collusion on the problem under consideration and we modify our algorithm to tolerate collusions. We conclude in Section VI. Omitted details and proofs can be found in the optional Appendix (to be read at the discretion of the program committee).

II. MODEL AND DEFINITIONS

We consider a distributed system consisting of n synchronous processes that can communicate via message-passing over a reliable network, where each process can communicate directly with each other process. Messages are not lost or corrupted in transit. Processes have unique ids from the set $[n] = \{1, \dots, n\}$.

The computation proceeds in synchronous rounds. In each round, each process can: (i) send point-to-point messages to selected processes, (ii) receive a set of point-to-point messages sent in the current round, and (iii) perform some local computation (if necessary). We assume that processes have access to a global clock, that is, rounds are globally numbered.

Processes may crash and restart dynamically as an execution proceeds. Each process is in one of two states: either *alive* or *crashed*. When a process is crashed, it does not perform any computation, nor does it send or receive any messages. We assume that processes have no durable storage, and thus when a process restarts, it is reset to a default initial state consisting only of the algorithm to execute and $[n]$. Each process can only crash or restart once per round. We denote by $crash(p, t)$ the event in which process p crashes in round t . The event $restart(p, t)$ is defined similarly. We say that a process p is *continuously alive* in the period $[t_a, t_b]$ if: (a) process p is *alive* at the beginning of round t_a and at the end of round t_b , and (b) for every $t \in [t_a, t_b]$, there are no $crash(p, t, \cdot)$ events.

When a process p crashes in round t , some of the messages sent by p in round t may be delivered, and some

may be lost. Similarly, when a process p restarts in round t , some of the messages sent to p may be delivered and some may be lost.

Rumors are dynamically injected into the system as the execution proceeds. A rumor ρ consists of a triplet $\langle z, d, D \rangle$, where z is the data to be disseminated, $D \subseteq [n]$ is the set of processes that z must be sent (destination set), and d is the deadline duration by which the rumor must be delivered. We denote by $Inj(\rho, t, p)$ the event in which rumor ρ is injected to process p in round t . We will be referring to p as the *source* process of rumor ρ . We assume that at most one rumor is injected at each process per round.

We model crash/restarts and rumor injection via a **Crash-and-Restart-Rumor-Injection adversary**, or **CRRI** adversary for short. In each round, the adversary determines which processes to fail, which processes to restart, and which rumors to inject. The adversary is *adaptive* in the sense that it can make decisions in a round t based on the events in all prior rounds $< t$, as well as the random choices being made in round t itself. We refer to an *adversarial pattern* $\mathcal{A} \in CRRI$ as a set of crash, restart and injection events caused by adversary **CRRI**.

Definition 1 (Quality of Delivery): We say that a gossip protocol guarantees *quality of delivery* if every rumor ρ injected in round t at a process p is delivered no later than round $t + \rho.d$ to every process in $\rho.D$ that is continuously alive for $[t, t+d]$, if p is also continuously alive for $[t, t+d]$.

Definition 2 (Confidentiality): We say that a gossip protocol is *confidential* if every rumor ρ is delivered only to processes in $\rho.D$, in every execution of the protocol.

Definition 3 (Per-round Message Complexity): We say that a *randomized* algorithm $Rand$ operating under adversary **CRRI** has *per-round message complexity* at most $M(Rand)$, if for every round t , for every $\mathcal{A} \in CRRI$, the following holds with high probability: the number of messages sent $M_t(Rand, \mathcal{A})$ by $Rand$ in round t , is at most $M(Rand)$.

Essentially, for randomized algorithms we require to guarantee Quality of Delivery (that is, admissible rumors are delivered with probability 1) and confidentiality with a probabilistic bound on the per-round message complexity. More on the rationale of Quality of Delivery and in general on the continuous gossip problem can be found in [12].

III. THE LIMITATIONS OF STRONG CONFIDENTIALITY

We say that a gossip protocol is *strongly confidential* if for every rumor, no message causally dependent on that rumor is ever sent to a process that is not in the destination set of that rumor. This essentially implies that only the processes in the destination set of a rumor can collaborate for that rumor's dissemination. As the following theorem states, such collaboration incurs high per-round message

complexity, even against an *oblivious adversary* that can only arrange the rumors destination sets prior to the start of the computation.

Theorem 1: For any constant $\varepsilon > 0$, every randomized strongly confidential gossip algorithm has a maximum per-round message-complexity of at least $\Omega(n^{(3/2)-\varepsilon}/dmax)$, with probability 1, even against an oblivious adversary, where $dmax$ is the longest deadline of the injected rumors.

Proof: (sketch) Let $x = n^{1/2-2/c}$ and let $c = \lceil 2/\varepsilon \rceil$. Suppose that only rumors with uniform deadlines $dmax$ are injected. We show a lower bound $\frac{nx}{2c \cdot dmax} = \Omega(nx/dmax) \supseteq \Omega(n^{(3/2)-\varepsilon}/dmax)$. Suppose that each process receives one rumor with random set of destinations in the beginning of the computation, defined independently over processes and in such a way that for each process it is decided independently with probability x/n whether it belongs to this destination set (or not, otherwise). The crucial argument is that under this scenario, with probability at least $1 - x^{2c+2}/n^{c-1}$, no message can carry more than c rumors. Having this, observe that the number of pairs (*source process, destination process*) is at least $nx/2$, with probability at least $1 - e^{-nx/8} \geq 1 - 1/e$, by the Chernoff bound. It follows that the total number of rumor copies carried out by messages is at least $nx/2$. Therefore, the total number of messages during delivering these rumors is at least $\frac{nx}{2c}$, with probability at least $1 - 1/e - x^{2c+2}/n^{c-1} \geq 1 - 1/e - n^{-2} \geq 1 - 2/e > 0$. By probabilistic argument, there is a configuration of destination sets that the above arguments are satisfied (i.e., at most c rumors per message, at least $nx/2$ pairs in communication schedule), and so the conclusion of at least $\frac{nx}{2c}$ messages to deliver all these rumors. This must be accomplished within $dmax$ rounds, thus there must be a round with at least $\frac{nx}{2c \cdot dmax}$ messages. Since we do not need crashes/restarts and the destination sets can be computed offline, the adversary is oblivious. ■

In view of the upper bound $O(n^{1+6\sqrt[3]{dmin}} \text{polylog } n)$ on continuous gossip without confidentiality assumptions [12] (against an adaptive adversary), we obtain a polynomial, in n , *price of strong confidentiality*, in terms of per-round message complexity (for sufficiently long deadlines, i.e., with minimum deadline $dmin > 24$). The above result has motivated our study of the weaker version of confidential gossip that allows processes outside a rumor's destination set to receive a message related to this rumor, as long as the rumor datum is not revealed.

IV. GOSSIPING CONTINUOUSLY AND CONFIDENTIALLY

In this section we present and analyze a continuous gossip algorithm, called CONGOS, that guarantees that the content of rumors remains confidential under adversary **CRRI**. For simplicity, we assume no collusion. In Section V, we show how to extend the approach here when processes collude. We

first describe the algorithm (Section IV-A) and then we show its correctness and analyze its performance (Section IV-B).

A. Algorithm CONGOS

In a nutshell, when a rumor is injected at a process p_i , the algorithm repeats the following procedure $\log n$ times concurrently: *Step 1*: Process p_i splits the rumor into two fragments such that only a process with both fragments can reconstruct the rumor. (In Section V-B, when processes collude, we will split each rumor into more fragments.) The processes are partitioned (deterministically) into two equal-sized groups. *Step 2*: Since process p_i itself belongs to one of the two groups, it uses a black-box continuous gossip service to share one of the half rumors with its own group. It uses a Proxy Service to distribute the other half rumor to the other group, with which it cannot gossip directly. At the end of the second step, each non-failed process has received one of the two half rumors. *Step 3*: The rumor fragments are sent to their appropriate final destinations using the GroupDistribution service. That is, the fragments for rumor ρ are sent to process in the destination set $\rho.D$. The algorithm guarantees that no process outside a rumor’s destination set gets both fragments of the rumor, while all processes in the rumor’s destination set (for which the rumor is admissible) deliver the rumor by the specified deadline. (A pseudocode-based description of the algorithm is given in Appendix A.)

We now proceed to present the technical details of the above outline.

1) *Preliminary Issues*: For the purposes of the algorithm description, we fix a deadline $dline$ and focus on rumors with deadlines in the range $[dline/2, dline]$. We then execute $\Theta(\log \log n)$ instances of the protocol, each for a specified range of deadlines. When given a rumor with some deadline $d \in [1, \Theta(\log^6 n)]$, we inject that rumor into the instance with the appropriate deadline range. When given a rumor with some deadline $> \Theta(\log^6 n)$, we truncate the deadline and inject it into the instance with the largest deadline. (There is no benefit to deadlines longer than $\Theta(\log^6 n)$.)

We present algorithm CONGOS as a set of composed distributed services. Each service is implemented by a protocol that is executed in a distributed fashion over all the processes in the system. This allows us to leverage existing distributed protocols as a black box, without delving into the underlying implementation details. We assume that the system consists of the following services:

- **Network**: We model the communication network as one such distributed service, with local input port *send* and a local output port *receive* at each process.
- **GroupGossip $[\ell]$** : We assume the availability of an existing *Continuous Gossip service*, albeit, one that does not guarantee confidentiality. It does ensure, however, that every rumor injected is delivered by the specified dead-

line, and it bounds the per-round message complexity by $O(n^{1+6\sqrt[3]{dmin}} \text{polylog } n)$, with high probability, where $dmin$ is the shortest deadline of any active rumor (see [12]).

We assume there are $\log n$ instantiations of this continuous gossip service, GroupGossip $[\ell]$ for $\ell \in \{1, \dots, \log n\}$. The instance GroupGossip $[\ell]$ is associate with partition ℓ of the network, which we define shortly. Every message sent by GroupGossip $[\ell]$ is *filtered* before being sent over the network: if a process p_i is a member of some group P' in partition ℓ , then every message sent by GroupGossip $[\ell]$ at process p_i to a process *not* in P' is dropped; every message sent by GroupGossip $[\ell]$ at process p_i to a process in P' is relayed to the Network and sent. From the perspective of the instance GroupGossip $[\ell]$, the processes that cannot be reached due to the filter are effectively failed.

- **AllGossip**: We assume a single continuous gossip service, AllGossip, that is not filtered. That is, it is allowed to communicate with all processes in the system.

The algorithm is composed into seven services, running in parallel at each process, and communicating via local ports. These are Network, GroupGossip, AllGossip, Filter, ConfidentialGossip, Proxy, and GroupDistribution. The ConfidentialGossip service is the main service that controls the flow of the algorithm and it coordinates the other services. We describe it next. (Figure 1 in Appendix A depicts the interaction of the various services at a process p_i .)

2) *Main Control: ConfidentialGossip Service*: Rumors are injected in the ConfidentialGossip service. In order to ensure confidentiality, each rumor ρ is divided into two fragments ρ_0 and ρ_1 . Both fragments maintain certain *metadata*, such as the rumor’s destination set, but each fragment on its own provides no information as to the original rumor datum $\rho.z$; yet together, they allow the original rumor to be reconstructed. There are a variety of simple schemes for accomplishing this: for example, let $\rho_0.z$ be a random binary string, and let $\rho_1.z = (\rho.z \mathbf{xor} \rho_0.z)$. In this way, we have reduced the problem of confidentiality to ensuring that no process, except those in the destination set, learn both fragments of the rumor. All the processes in the system are partitioned into two components, and one rumor fragment is distributed to each half.

It is not sufficient, however, to carry out this splitting-and-partitioning process only once: the adversary, being adaptive, may kill all the processes in one of the groups in the partition. We thus define, a priori, $\log n$ different partitions. Each partition is based on a specified bit in the binary representation of a process’s identifier. Let $p_j[\ell]$ be the ℓ^{th} bit in p_j ’s binary representation. Then partition ℓ is defined by the two sets $P_{0,\ell} = \{p_j : p_j[\ell] = 0\}$ and $P_{1,\ell} = \{p_j : p_j[\ell] = 1\}$. For each partition, rumor ρ is divided into two fragments $\rho_{0,\ell}$ and $\rho_{1,\ell}$.

Therefore, the ConfidentialGossip service, running at each

Outline of ConfidentialGossip service at p_i :

- Do in parallel for each $\ell = 1, \dots, \log n$:
 - 1) Split rumor ρ into a pair $\langle \rho_{0,\ell}, \rho_{1,\ell} \rangle$.
 - 2) If p_i is in group b of partition ℓ , inject $\rho_{b,\ell}$ into GroupGossip $[\ell]$, and inject $\rho_{1-b,\ell}$ into Proxy $[\ell]$. Together, these two services ensure that each rumor fragment is delivered to every non-failed process in the appropriate group of the partition.
 - 3) For each rumor fragment received from GroupGossip $[\ell]$ or Proxy $[\ell]$, inject the fragment into GroupDistribution $[\ell]$.
 - 4) Save every fragment received from GroupDistribution $[\ell]$, and reassemble and deliver rumors as fragments become available.
 - Whenever a message from AllGossip confirms that, for some partition ℓ , both fragments of a rumor ρ , initiated at p_i , have been sent to every destination in $\rho.D$, confirm that ρ has been delivered.
 - Whenever a deadline is about to expire for some rumor ρ initiated at p_i , and there is no confirmation that ρ has been delivered, send ρ directly to every process in $\rho.D$.
-

process p_i , “spawns” ℓ instances of the other services. We will be using the notation ServiceName $[\ell]$ to denote the instance of a service for partition ℓ . (The service AllGossip runs one instance for all partitions, as its purpose is to gossip information to all processes. This is also the case for the Network service.) Above is a high-level outline of the ConfidentialGossip at a process p_i . (Detailed pseudocode is given in Figure 2 in Appendix A.)

Time is divided into blocks of $dline/4$ rounds. A rumor injected during some block B is split into fragments during block B (step 1, above); the fragments are distributed to their respective groups during block $B + 1$ (step 2, above); the fragments are reassembled in block $B + 2$ (step 3 and 4, above); and the source verifies that its rumor has been delivered during block $B+3$. If it cannot verify that its rumor has been delivered when the deadline expires, it simply sends its rumor directly (last bullet above).

A notable aspect of the above protocol is that a process cannot directly distribute both rumors. If a process p_i is in group $P_{0,\ell}$, it cannot directly participate in gossip with group $P_{1,\ell}$; if it did, it might risk learning rumor fragments associated with the other group. The Proxy service is designed to circumvent this problem.

Another notable aspect occurs at the end of the protocol, when a process attempts to confirm that its rumors have been delivered. Each process, as part of the GroupDistribution service (see below), initiates a gossip (via AllGossip) indicating which rumor fragments have been distributed to which processes. Of course, a process cannot divulge the *contents* of the rumor that have been distributed; however, it can safely indicate a unique identifier that was appended by the source, when the rumor was split. In this way, the source can ensure that, for at least one partition, both rumor fragments were successfully delivered. (Note that it would not be sufficient for recipients to send an acknowledgment,

as the source does not know which processes have remained alive throughout the interval.) We now outline the Proxy service and the GroupDistribution service.

3) *Proxy Service*: The goal of the proxy service is to deliver rumor fragments safely across group boundaries. Essentially, the proxy service for partition ℓ at process p_i repeatedly samples processes from the other group (i.e., the group that p_i does not belong to), requesting that these processes act as *proxy* for p_i in distributing its rumor fragments. The potential proxies then participate in GroupGossip $[\ell]$, attempting to distribute the rumor fragments, as requested. If they succeed, they send an acknowledgment to p_i . Otherwise, process p_i needs to try again.

The challenge, here, is that the adversary may (adaptively) crash processes as soon as they receive proxy requests. (In fact, for some partitions, the adversary may crash all the members of a given group.) Even worse, at any given time, most of the members of a group may be failed, requiring p_i to send a very large number of queries to find a proxy. To avoid this problem, the processes in the same group collaborate on finding proxies. At the same time, p_i does not share any information on the fragments it is attempting to distribute in the *other group* with processes in the *same group* with which it is collaborating. The Proxy service for partition ℓ proceeds as follows (detailed pseudocode is given in Figure 3 in Appendix A):

Outline of Proxy $[\ell]$ at p_i :

- Time is divided into blocks of length $dline/4$.
 - At the beginning of a block, collect all the fragments that have been injected since the last block began, and set status to active.
 - Each block is divided into *iterations* of $\sqrt{dline} + 2$ rounds. In each iteration, we maintain a set *collaborators* of the active processes in the same group as p_i . We also keep track of *failed-proxies*, i.e., those that we have already learned (in previous iterations) have failed in this block. For each iteration, repeat:
 - Round 1: send every rumor fragment associated with the other group to $n^{1+48/\sqrt{dline}} \log n / |\text{collaborators}|$ processes chosen uniformly at random from the other group, excluding processes in *failed-proxies*. (Notice that as long as the set *collaborators* is a good estimate of the set of collaborators, this ensures a good bound on the message complexity of this step.) Every process that receives a request to be a proxy for the other group caches the received rumor fragments.
 - Rounds 2, \dots , $\sqrt{dline} + 1$: initiate a GroupGossip $[\ell]$ in which processes in the same group as p_i share the set of *failed-proxies*, as well as establish the set of *collaborators*, i.e., members of the group that are still active. Processes also share all the rumor fragments received from the other group. (The deadline for rumors in GroupGossip $[\ell]$ here is \sqrt{dline} .)
 - Round $\sqrt{dline} + 2$: Any process that was asked to be a proxy for the other group sends an acknowledgment that proxying was successful. Any process that sent a request, and does not receive an acknowledgment, adds the non-acknowledging processes to the set of *failed-proxies*.
-

4) *GroupDistribution Service*: The goal of the $\text{GroupDistribution}[\ell]$ service is to distribute rumor fragments to their final destination. To this point, for a partition ℓ , the rumor fragment $\rho_{0,\ell}$ has been distributed to processes in $P_{0,\ell}$, and the rumor fragment $\rho_{1,\ell}$ has been distributed to processes in $P_{1,\ell}$. Now, group $P_{0,\ell}$ collaborates to send the fragment $\rho_{0,\ell}$ to $\rho_{0,\ell} \cdot D$, while $P_{1,\ell}$ does the same for fragment $\rho_{1,\ell}$. Of course there may be many different fragments active in each group, each with a different destination set.

The basic operation of the GroupDistribution is similar to that of the Proxy Service . Each process chooses a set of recipients at random, and sends each of them a message carefully composed to only include appropriate rumor fragments. The processes then gossip within their group (via $\text{GroupGossip}[\ell]$), sharing information on which processes have already been notified, and which remain to be notified. At the same time, processes calculate the number of processes active in a group, which allows them to determine the appropriate number of messages to send. Below we give an outline of the GroupDistribution service for partition ℓ (detailed pseudocode is given in Figure 4 in Appendix A.)

B. Algorithm Analysis

We begin the analysis of algorithm CONGOS by stating its correctness (i.e., confidentiality is not violated and all admissible rumors are delivered on time). Due to space limitation, its formal proof is deferred to Appendix B.

Theorem 2 (Correctness): Algorithm CONGOS correctly solves the Confidential Continuous Gossip problem under adversary CRRI .

In the next few lemmas, we state important properties needed for analyzing the message complexity of the algorithm. Full or omitted proofs can be found in Appendix B.

Lemma 3: Given rumor ρ , injected at time t : if there are at least 2 processes that remain alive throughout the interval $[t, t + \rho \cdot d]$, then for some partition ℓ , there is at least one process in $P_{0,\ell}$ and one process in $P_{1,\ell}$ that remain alive throughout the interval $[t, t + \rho \cdot d]$.

Proof: Let p_i and p_j be the two processes hypothesized to remain alive throughout the specified interval. Since identifiers are unique, let ℓ be some bit where the identifier of p_i and p_j differ. The claim follows for partition ℓ . ■

Lemma 4: In each block, the $\text{Proxy}[\ell]$ service and the $\text{GroupDistribution}[\ell]$ service execute at least $\sqrt{dline}/8$ iterations, if $dline > 4$.

Proof: Each block is of length $dline/4$, and each interval is of length at most $\sqrt{dline} + 2 \leq 2\sqrt{dline}$. ■

Lemma 5: In each round, the $\text{Proxy}[\ell]$ service and the $\text{GroupDistribution}[\ell]$ service send at most $O(n^{1+48/\sqrt{dline}} \log n)$ messages.

Outline of $\text{GroupDistribution}[\ell]$ at p_i :

- Time is divided into blocks of length $dline/4$.
 - At the beginning of the second round of a block, collect all the fragments that have been injected since the first round of the block, and set status to active. (The first round of the block is spent waiting for rumor fragments from the previous block.)
 - Each block is divided into *iterations* of $\sqrt{dline} + 2$ rounds. In each iteration, we maintain a set *collaborators* of the active processes in the same group as p_i . We also keep track of a set *hitSet* of processes that have been sent a message in this block. Each process in this set was sent all the rumor fragments for this block. For each iteration, repeat:
 - Round 1: wait for rumor fragments to be injected.
 - Round 2: send every “appropriate” rumor fragment to $n^{1+48/\sqrt{dline}} \log n / |\text{collaborators}|$ processes chosen uniformly at random from the other group, excluding processes in *hitSet*. By appropriate we mean that if p_j is a process chosen randomly by p_i , then p_i sends to p_j only the rumor fragments in which p_j is in the destination set. (Recall that each partial rumor contains the target destination set as part of the metadata.) Every process that receives rumor fragments can now reconstruct the rumor and return it to its user (via the $\text{ConfidentialGossip}$ service).
 - Rounds 3, \dots , $\sqrt{dline} + 2$ rounds: initiate an instance of $\text{GroupGossip}[\ell]$ (with deadline \sqrt{dline}) in which processes in the same group as p_i share their *hitSets*, as well as count how many members of the group are still active.
 - In the last round of the block, initiate an instance of AllGossip (with deadline $dline/4 - 1$). Each process p_i gossips the information in its *hitSet*, but without including the rumor fragments themselves. That is, if the *hitSet* of process p_i indicates that some rumor fragment $\rho_{0,\ell}$ was sent to some process p_j , and if $\rho_{0,\ell}$ has identifier r , then p_i gossips that the fragment 0 for partition ℓ of the rumor associated with identifier r was sent to p_j . This provides sufficient information for the source to determine whether the rumor was delivered, without revealing the contents of the rumor. (See the description of the $\text{ConfidentialGossip}$ service, above, for how this information is used.)
-

Proof: (sketch) In $\text{Proxy}[\ell]$ and $\text{GroupDistribution}[\ell]$, each process sends $\frac{n^{1+48/\sqrt{dline}} \log n}{|\text{collaborators}|}$ messages in, respectively, the first and second round of an iteration. The bound on *collaborators* implies the desired result. In $\text{Proxy}[\ell]$, each process that received a proxy request sends a response at the end of an iteration. Each response is the result of an earlier request in the first round of the iteration, and we have already bounded the message complexity of the first round of an iteration, leading here too to a bound of $O(n^{1+48/\sqrt{dline}} \log n)$. ■

The challenge of showing the next lemma lies on the fact that the adversary is adaptive. If the adversary were oblivious, then we could simply analyze the random choices as independent. But because the adversary is adaptive and can schedule according to the random choices, there is subtle correlation among the random choices, and hence we have

to show that the requisite properties hold despite all possible adversarial choices.

Lemma 6: Given rumor ρ , injected at time t at process p_i : if at least one process in $P_{0,\ell}$ and one process in $P_{1,\ell}$ remain alive throughout the interval $[t, t + \rho.d]$, then every process in $P_{0,\ell}$ that remains alive throughout the interval $[t, t + \rho.d]$ receives $\rho_{0,\ell}$, and every process in $P_{1,\ell}$ that remains alive throughout the interval $[t, t + \rho.d]$ receives $\rho_{1,\ell}$ by time $t + 2dline/4 - (t \bmod dline)$, with high probability.

Proof: (sketch) Fix ℓ to be the partition identified in Lemma 3. Assume w.l.o.g. that $p_i \in P_{0,\ell}$. (The alternate case is symmetric.) Since p_i injects rumor fragment $\rho_{0,\ell}$ into the GroupGossip $[\ell]$ service with deadline \sqrt{dline} , it is guaranteed to reach every process in $P_{0,\ell}$ that remains alive throughout the interval. It remains to show that each process in $P_{0,\ell}$ succeeds in finding a proxy in $P_{1,\ell}$, while executing Proxy $[\ell]$ during the first complete block after rumor ρ is injected beginning at time $t + dline/4 - (t \bmod dline)$. Let $Z = n^{48/\sqrt{dline}}$. The crucial argument is that in every pair of iterations, one of the three following events occurs: (i) At least a $(1 - 1/Z)$ fraction of processes in $P_{0,\ell}$ that were alive at the beginning of the first iteration fail by the end of the second iteration; (ii) At least a $(1 - 1/Z)$ fraction of processes in $P_{1,\ell}$ that were alive at the beginning of the first iteration fail by the end of the second iteration; (iii) In the second iteration, at least a $(1 - 1/Z)$ fraction of processes in $P_{0,\ell}$ succeed in finding a proxy. From this claim and by the lemma assumptions, we can conclude that by the end of $3 \log_Z(n)$ pairs of iterations, every process in $P_{0,\ell}$ has succeeded in finding a proxy. Note that since $\log_Z(n) = \sqrt{dline}/48$, this process finishes within $6 \log_Z(n) \leq \sqrt{dline}/8$ iterations, as required. Once a process has succeeded in finding a proxy, it follows from the guarantees of GroupGossip that its rumor fragment is distributed to every non-failed process in the other group. ■

Lemma 7: Given rumor ρ , injected at time t at process p_i : if at least one process in $P_{0,\ell}$ and one process in $P_{1,\ell}$ remain alive throughout the interval $[t, t + \rho.d]$, then every process $p_j \in \rho.D$ receives fragment $\rho_{0,\ell}$ and fragment $\rho_{1,\ell}$ by time $t + 3dline/4 - (t \bmod dline)$, with high probability.

Proof: (sketch) Due to Lemma 6, it remains to show that during the following block of rounds, for every process $p_j \in \rho.D$, at least one process from each group sends its rumor fragment to p_j . W.l.o.g. we focus on group $P_{0,\ell}$. Let $Z = n^{48/\sqrt{dline}}$. The crucial observation is that in each pair of iterations of the GroupDistribution $[\ell]$ service, one of the following two events occurs: (i) At least a $(1 - 1/Z)$ fraction of processes in $P_{0,\ell}$ that were alive at the beginning of the first iteration fail by the end of the second iteration; (ii) For every process p_k active throughout both iterations, the set of processes $[n] \setminus hitProcs_k$ decreases by a factor of Z by the end of the second iteration. ($hitProcs_k = \{p_q \in [n] :$

$\langle p_q, \cdot \rangle \in hitSet_k\}$.) From this claim, we conclude that within $2 \log_Z n$ pairs of iterations, either every process in $P_{0,\ell}$ fails, or every process has been added to $hitProcs$. Hence we conclude that by the end of $4 \log_Z n \leq \sqrt{dline}/8$ iterations, every process has been sent all of its rumor fragments. ■

Lemma 8: Given rumor ρ , injected at time t at process p_i : if p_i does not fail by time $t + \rho.d$, then by round $t + \rho.d - 1$, process p_i receives confirmation that rumor ρ was delivered, with high probability.

Proof: By Lemma 3, we know that if rumor ρ has even one admissible destination $\neq p_i$, then there is some partition ℓ where there is at least one process in $P_{0,\ell}$ and one process in $P_{1,\ell}$ that does not fail in $[t, t + \rho.d]$. By Lemma 7, we know that by time $t + 3dline/4 - (t \bmod dline)$, with high probability, rumor ρ has been delivered to every destination in $\rho.D$ by the GroupDistribution $[\ell]$ service. Moreover, since at least one process p_0 in $P_{0,\ell}$ and one process p_1 in $P_{1,\ell}$ does not fail during $[t, t + \rho.d]$, we conclude that p_0 and p_1 complete the block in which the rumor fragments for ρ are delivered to their destinations. At the end of the last round of the block, processes p_0 and p_1 inject sanitized versions of the $hitSets$ as rumors into the AllGossip service with deadline $dline/4 - 1$, thus ensuring that process p_i receives this information no later than round $t + \rho.d - 1$. Process p_i then marks rumor ρ confirmed. ■

Theorem 9 (per-round message complexity): The per-round message complexity of CONGOS is:

$$O\left((n^{1+48/\sqrt{dline}} + n^{1+6/\sqrt[6]{dline}}) \text{polylog } n\right).$$

Proof: From Lemma 8 we have that for a given rumor ρ injected at process p_i , with high probability the rumor is confirmed prior to the deadline expiring. Since each process is injected at most one rumor per round (hence there can be $O(n \text{polylog } n)$ active rumors in the system at any given time), with high probability, no source process sends any messages directly to the destinations. From Lemma 5, the per-round message complexity for each instance of Proxy $[\ell]$ and GroupDistribution $[\ell]$ is $O(n^{1+48/\sqrt{dline}} \log n)$, leading to a per-round message complexity of $O(n^{1+48/\sqrt{dline}} \log^2 n)$. Each instance of continuous gossip, invoked with rumors at least \sqrt{dline} , has message complexity $O(n^{1+6/\sqrt[6]{dline}} \text{polylog } n)$. There are $\log n + 1$ such instances of continuous gossip. ■

V. GOSSIPING IN THE PRESENCE OF COLLUSION

In this section we extend our investigation of the confidential gossip problem by additionally assuming that processes outside of a rumor's destination set may collude in an attempt to learn the rumor.

More formally, given a rumor ρ injected in the system at a process p_i , we denote by C_ρ the collusion set of ρ . In particular, C_ρ may contain any process $q \notin \rho.D \cup \{p_i\}$.

We assign adversary $CRR I$ with the additional task of “selecting” the colluding processes in an adaptive way during the execution. We will be referring as $CRR I(\tau)$ the subset of the adversarial patterns of $CRR I$ for which $|C_\rho| \leq \tau$ for any rumor ρ injected in the system. Finally, we will be calling τ -collusion-tolerant an algorithm that it is designed to solve confidential gossip under adversary $CRR I(\tau)$.

A. Lower Bound

We prove that a class of algorithms generalizing our algorithm CONGOS suffers from collusion, in terms of message complexity. We say that a gossip algorithm is *partition-based* if it allows only two operations tampering with the content of the rumors: splitting, which allows to split a given initial rumor into disjoint smaller fragments, and merging, which allows to merge given fragments of the *same* rumor into a larger fragment of this rumor¹. Otherwise, the protocol must treat the rumor (and its fragments) as nonmalleable tokens.

The effect of collusion, as demonstrated by the following theorem, might be significant (especially for large number of colluders), even against an *oblivious adversary* that can only arrange the rumors destination sets and identify the colluding processes prior to the start of the computation.

Theorem 10: For any constant $\varepsilon > 0$, every randomized, τ -collusion-tolerant, partition-based algorithm solving confidential gossip has a maximum per-round message-complexity of at least $\Omega(\min\{n\tau, n^{(3/2)-\varepsilon}\}/dmax)$, with probability 1, against an oblivious adversary, where $dmax$ is the longest deadline of the injected rumors.

Proof: We assume the same initial setting of parameters and rumors, including their destination sets and deadlines, as considered in the proof of Theorem 1, which are as follows. We may assume that n is sufficiently large (in fact, $n \geq 8$ is sufficient). Let c be a constant and x be a parameter, to be specified in the same way as in the proof of Theorem 1, depending on ε . Suppose that only rumors with uniform deadlines $dmax$ are injected, all at the same time. Moreover, assume that each process is injected one rumor with the same destination set as in the proof of Theorem 1.

Now consider a single execution of a given algorithm in this setting. Let a *rumor interval* be a set of rumor fragments such that any set of fragments sufficient to reconstruct the rumor includes one fragment from the rumor interval. (Informally a rumor interval corresponds to a sub-sequence of rumor bit-string representation and to all rumor fragments that contain this sub-sequence.) Two cases are possible:

Case 1: More than half of the rumors satisfy the following property each: there is a rumor interval such that none of its

contained fragments is ever transmitted to a process outside the destination set.

It follows that such a rumor interval, each destination process receives some rumor fragment in this rumor interval directly from the rumor’s source (the process that the rumor was injected at) or relayed entirely through the processes in the destination set. Therefore, the messages carrying rumor fragments in this rumor interval altogether suffer from the same constraints as it would an original rumor propagated within its destination sets only. By Theorem 1, the number of such messages is proportional to the size of the destination set, for the considered setting of destination sets, and since there are more than $n/2$ such rumors, we get the lower bound $\Omega(n^{(3/2)-\varepsilon})$ on the total number of such messages in the considered period of length $dmax$. Hence, the per round message complexity in this case is $\Omega(n^{(3/2)-\varepsilon}/dmax)$.

Case 2: At least half of the rumors satisfy the following property each: fragments of the rumor transmitted outside the destination set cover the whole original rumor.

In this case for each such rumor there are at least $\tau + 1$ processes outside its destination set that receive a fragment of the rumor directly from some processes in the destination set or the rumor’s source (the process that the rumor was injected at); otherwise at most τ such outside processes could collude and get fragments covering the whole rumor, thus violating the definition of confidentiality (which must hold for every execution). We call such at least $\tau + 1$ fragments *border fragments*. Therefore there are at least $\tau + 1$ point-to-point messages sent from some processes in the destination set or the rumor’s source to the considered at least $\tau + 1$ outside processes. Call these messages *border messages*. It follows that there are at least $(\tau + 1)n/2$ copies of border fragments sent via border messages. Recall the property of the considered configuration of destination sets as proved in Theorem 1: each process is in at most c destination sets, where c is a constant. It follows that a process sends at most c border fragments per border message, which gives at least $\frac{(\tau+1)n/2}{c} = \Omega(n\tau)$ border messages. Hence the per-round message complexity in this case is $\Omega(n\tau/dmax)$.

In each case, the message complexity is $\Omega(\min\{n\tau, n^{(3/2)-\varepsilon}\}/dmax)$. This bound holds for any execution of the algorithm. The adversary is oblivious, as it uses the same setting as in the proof of Theorem 1 and does not need to specify colluding processes. To justify the latter, observe that both cases hold regardless of the choice of colluding processes, and the only place the adversary threads the algorithm by possibility of collusion is in Case 2 when it enforces at least $\tau + 1$ border messages; but for this it does not need to specify online the set of colluding processes, and the argument says only that if the algorithm broke it, the adversary could choose a set of colluders violating confidentiality. ■

¹Notice that this does not allow other algebraic manipulation of the rumor, as in “network coding” techniques.

B. Algorithm

We modify algorithm CONGOS in the following way. Instead of $\log n$ partitions used in algorithm CONGOS, we use $c\tau \log n$ partitions given as a part of the input of the algorithm, for an appropriate choice of constant c . Each partition contains $\tau+1$ groups, instead of the originally used 2 groups. For this purpose, rumors are now divided into $\tau+1$ fragments. If we view $\tau = 1$ as a collusion of a process with itself, then the original algorithm CONGOS can be viewed as 1-collusion-tolerant confidential gossip algorithm.

The set of $c\tau \log n$ partitions needs to satisfy the following properties, for appropriate choice of constants c and c' :

- In each partition, each group contains at least one process.
- For every set S of at least $2c'\tau \log n$ processes, there exists a partition such that every group in the partition contains at least one process in S .

The first property ensures well-formedness, i.e., that the partition is a proper division of the processes into non-empty groups. The second property ensures good performance: as long as there are $\Omega(\tau \log n)$ processes alive, then one partition has live processes in every group and hence can be used to distribute the rumor fragments. We now argue that there exists a good set of partitions that meets these requirements:

Lemma 11: If $\tau < n/\log^2 n$, then there is a set of $c\tau \log n$ partitions satisfying the above conditions, for some constants $c, c' > 0$.

Proof: We proceed via the probabilistic method: first, we randomly select the groups for each partition, and then we show that with some positive probability, the two properties are satisfied. We begin by assuming that for each partition, each process is independently assigned, uniformly at random, to one of the $\tau+1$ groups in that partition.

We begin by examining the first required property. For each group g , the probability that a process chooses group g is $1/(\tau+1)$. Thus the probability that a given group is not chosen by any process is at most $(1-1/(\tau+1))^n \leq 2^{-\log^2 n}$. In total, there are $(\tau+1) \cdot (c\tau \log n) \leq 2cn^2$ groups, and thus by a union bound, the probability that any group is not chosen by some process is at most $2^{-(\log^2 n - \log(2cn^2))} < 1/2$, for sufficiently large n .

We now examine the second required property. We fix some set S of size $2c'\tau \log n$. For a given partition, for a given group in that partition, the probability that no process in set S is assigned to that partition is at most $(1-1/(\tau+1))^{2c'\tau \log n} \leq 1/n^{c'}$. Thus, by a union bound, the probability that any group in the partition does not contain a process in S is at most $(\tau+1)/n^{c'} \leq 1/n^{c'-1}$.

Since each partition is selected independently, the probability that for every one of the $c\tau \log n$ partitions, at least one group is empty is at most $(1/n^{c'-1})^{c\tau \log n} \leq n^{-c \cdot c' \tau \log n/2}$.

Now, consider all $\binom{n}{2^{c'\tau \log n}}$ choices of the set S . There are at most $n^{2c'\tau \log n}$ such sets S . Taking a union bound over all the sets S , the probability that there exists a set S for which every partition has at least one empty group is at most $n^{-(c \cdot c' \tau \log n/2 - 2c'\tau \log n)} < 1/2$, for appropriately large n and choice of c and c' .

Thus, the probability that the selected partition does not satisfy the two requisite properties is smaller than 1, and hence, by the probabilistic method, a partition satisfying the two desired properties exists. ■

We leave the polynomial time construction of partitions satisfying the required conditions as future work.

Overview of collusion-tolerant CONGOS: In a nutshell, the modified version of algorithm CONGOS operates as follows for a newly injected rumor ρ at a process p_i . Procedure ConfidentialGossip is called in which the rumor, for each different partition ℓ , is divided into the fragments $\rho_{0,\ell}, \dots, \rho_{\tau,\ell}$ such that all fragments (from the same partition) are needed in order for ρ to be re-assembled. (A way to do this is as follows: Let $\rho_{0,\ell}, \dots, \rho_{\tau-1,\ell}$ be different random binary strings and set $\rho_{\tau,\ell} = (\rho \mathbf{xor} \rho_{0,\ell} \mathbf{xor} \dots \mathbf{xor} \rho_{\tau-1,\ell})$. Then ρ can be computed when all $\tau+1$ fragments are received. Note that this scheme makes the algorithm partitioned-based.)

Say that in partition ℓ , process p_i belongs in group x . Then it injects fragment $\rho_{x,\ell}$ in GroupGossip $[\ell]$ and all other fragments into Proxy $[\ell]$. Via procedure GroupGossip $[\ell]$, the fragment $\rho_{x,\ell}$ is gossiped in the members of group x and via Proxy $[\ell]$ each other fragment is gossiped into every other corresponding group (such that every other group learns a different fragment of the rumor). Then procedure GroupDistribution $[\ell]$ is called so that the processes in each group collaborate in sending their corresponding fragment of the rumor only to the processes of the rumor's destination set. These processes receive all fragments and hence can reassemble the rumor. Lemma 11 assures the existence of at least one partition ℓ in which all admissible rumors are received by the live processes of the rumor's destination set. Detailed outlines of the procedures are given in Appendix C.

We now give the main result of this section.

Theorem 12: The modified version of algorithm CONGOS solves the confidential gossip problem under adversary $CRR1(\tau)$ with per-round message complexity of

$$O\left(\left(n^{1+48/\sqrt{dline}} + n^{1+6/\sqrt[6]{dline}}\right)\tau^2 \text{polylog } n\right).$$

Proof: (sketch) The correctness follows by similar arguments as for algorithm CONGOS, since the partitions used for the modified algorithm (with the properties proved in Lemma 11) satisfy the same conditions explored in the analysis as the partitions used in the original algorithm CONGOS.

The message complexity increases by a factor τ^2 , compared to the complexity of the original algorithm CONGOS,

because of the following two observations.

First, for all rounds but the last one of the considered deadline period the amount of inter-group communication (that is, the last round in procedure GroupDistribution) increases by a factor of at most $\tau + 1$, as the number of groups is now $\tau + 1$ instead of 2. The communication is increased by another factor $\Theta(\tau)$ due to the fact that now there are $\Theta(\tau \log n)$ partitions instead of $\log n$. Thus each message sent in the original algorithm CONGOS is multiplied by at most $\Theta(\tau^2)$ different copies.

Second, in the last round of ConfidentialGossip (shooting directly to processes in the destination set) there is no communication if the number of processes alive in the whole period is at least $2c'\tau \log n$ (that is, all rumors have been confirmed to be delivered); this follows by the second property of the set of partitions and by the same argument as in the analysis of the original algorithm CONGOS. In the case where the number of alive processes is smaller than $2c'\tau \log n$, the number of point-to-point messages sent is bounded by $2c'\tau n \log n$ due to the fact that only these processes that remain alive throughout may send messages in the last round (each sending at most n messages). This completes the proof. ■

Observe from Theorem 12 that when $dline = \Theta(\log^6 n)$ the per-round message complexity is $O(n\tau^2 \text{polylog } n)$. When contrasted with Theorem 10 it follows that for $\tau < n^{1/4}$, the per-round message complexity is within a factor of $\tau \text{polylog } n$ of the lower bound. (For $\tau = O(\text{polylog } n)$ the algorithm is optimal within log factors.)

VI. DISCUSSION

In this paper we have considered the problem of *confidential* gossip, where each rumor is learned only by processes in the rumor’s specified destination set. Assuming an adaptive and omniscient adversary that dynamically and continuously injects rumors into the system and causes process crashes and restarts, we have designed an efficient (w.r.t. per-round message complexity) algorithm which we call algorithm CONGOS. As an alternative to cryptographic schemes, which can be expensive in such a dynamic environment, the algorithm deploys a simple rumor splitting technique that enables an efficient “all-process” collaboration while guaranteeing confidentiality. For this purpose, the algorithm combines, in a non-trivial way, a black-box efficient non-confidential continuous gossip service with other auxiliary services (namely, Filter, Proxy, GroupDistribution). While we have focused on continuous gossip, we believe that the same techniques apply to other gossip variants (e.g., single-instance gossip, etc.).

We have also discussed the problem of collusion, and shown how to tolerate a moderate amount of collusion at a limited cost. An interesting open question is whether we can tolerate higher levels of collusion if the adversary is

oblivious, or if we allow some small probabilistic violation of confidentiality.

In addition, as currently presented, the algorithm guarantees the confidentiality of rumors, but various other metadata is released. For example, processes learn of the existence of rumors, roughly how many rumors are active, the source of each rumor, a sequence number of each rumor, and the set of destinations for each rumor. Some of this information can be readily hidden. For example, the sequence number can be replaced with a pseudorandom identifier. Other information appears more difficult to hide, for example, the proxies learn precisely who is requesting that they act as a proxy, and this seems, to some extent, unavoidable.

The destination set associated with each rumor can be hidden, without increasing the overall message complexity, but at the cost of increasing the message size (significantly). When a rumor ρ is injected at process p_i , the source creates n new rumors, each with a single process in its destination set. For every process in $\rho.D$, the new rumor contains a copy of the injected rumor’s content. For the remaining new rumors, the contents of the new rumor are chosen at random. The source then proceeds to distribute this entire collection of rumors. Only the processes in the destination set can determine whether a rumor contains real content or simply a random string, and hence processes cannot determine the real destination set.

Similarly, the very existence of rumors can be hidden by the continual injection of fake content-free rumors, at the cost of wasted messages. In this way, a process cannot determine how many real rumors are currently active.

Finally, an interesting open question is whether we can tolerate truly malicious processes, i.e., those that do not follow the protocol. In fact, we believe that the approach for tolerating collusion may be extended to deal with malicious processes, if the adversary is oblivious. In that case, we can tolerate some groups misbehaving and failing to deliver their message fragments.

REFERENCES

- [1] S. Baehni, P.T. Eugster, and R. Guerraoui. Data-Aware Multicast. In *DSN 2004*, pages 233–242.
- [2] A.J. Ballardie. *A New Approach to Multicast Communication in a Datagram Network*, Ph.D. Thesis, University College London, 1995.
- [3] A. Beimel, K. Nissim, and E. Omri. Distributed Private Data Analysis. In *CRYPTO 2008*, pages 451–468.
- [4] J. Brickell and V. Shmatikov. Privacy-preserving Graph Algorithms in the Semi-honest Model. In *ASIACRYPT 2005*, pages 236–252.
- [5] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. Multicast Security: A Taxonomy and Some Efficient Constructions. In *INFOCOM 1999*, pages 708–716.
- [6] B.S. Chlebus and D.R. Kowalski. Time and Communication Efficient Consensus for Crash Failures. In *DISC 2006*, pages 314–328.

- [7] G. Chockler, R. Melamed, Y. Tock, and R. Vitenberg. SpiderCast: A Scalable Interest-Aware Overlay for Topic-Based Pub/Sub Communication. In *DEBS 2007*, pages 14–25.
- [8] G. Chockler, R. Melamed, Y. Tock, and R. Vitenberg. Constructing Scalable Overlays for Pub/Sub with Many Topics. In *PODC 2007*, pages 109–118.
- [9] G. Delpeste-Gallet, H. Fauconnier, R. Guerraoui, and E. Ruppert. Secretive Birds: Privacy in Population Protocols. In *OPDIS 2007*, pages 329–342.
- [10] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom Rumor Spreading: Expanders, Push vs Pull, and Robustness. In *ICALP 2009*, pages 366–377.
- [11] A. Fiat and M. Naor. Broadcast Encryption. In *CRYPTO 1993*, pages 480–491.
- [12] Ch. Georgiou, S. Gilbert, and D.R. Kowalski. Meeting the Deadline: On the Complexity of Fault-Tolerant Continuous Gossip. In *PODC 2010*, pages 247–256.
- [13] O. Goldreich. *Foundations of Cryptography: Volume II (Basic Applications)*. Cambridge University Press, 2004.
- [14] I. Gupta, A.M. Kermarrec, and A.J. Ganesh. Efficient Epidemic-style Protocols for Reliable and Scalable Multicast. In *SRDS 2002*, pages 180–189.
- [15] Havard D. Johansen, Andre Allavena, and Robbert van Renesse. Fireflies: Scalable Support for Intrusion-tolerant Network Overlays. In *EuroSys 2006*, pages 3–13.
- [16] J. Hromkovic, R. Klasing, A. Pelc, P. Ruzika, and W. Unger. *Dissemination of Information in Communication Networks: Broadcasting, Gossiping, Leader Election, and Fault-Tolerance*, Springer-Verlag, 2005.
- [17] R. Karp, C. Schindelhauer, S. Shenker, B. Vocking. Randomized Rumor Spreading. In *FOCS 2000*, pages 565–574.
- [18] D. Kempe, J. Kleinberg, and A. Demers. Spatial Gossip and Resource Location Protocols. *Journal of the ACM*, 51:943–967, 2004.
- [19] A. Kermarrec, L. Massoulie, A. Ganesh. Probabilistic Reliable Dissemination in Large-scale Systems. *IEEE Transactions on Parallel and Distributed Systems*, 14(3):248–258, 2003.
- [20] L. Kissner and D. Song. Privacy-preserving Set Operations. In *CRYPTO 2005*, pages 241–257.
- [21] D. R. Kowalski and M. Strojnowski. On the Communication Surplus Incurred by Faulty Processors. In *DISC 2007*, pages 328–342.
- [22] Y. Lindell and B. Pinkas. Privacy Preserving Data Mining. *J. Cryptology*, 15(3):177–206, 2002.
- [23] D. Malkhi, Y. Mansour, and M.K. Reiter. Diffusion Without False Rumors: On Propagating Updates in a Byzantine Environment. *Theoretical Computer Science*, 299:289–306, 2003.
- [24] D. Micciancio and S. Panjwani. Corrupting One Vs. Corrupting Many: The Case of Broadcast and Multicast Encryption. In *ICALP 2006*, pages 70–82.
- [25] Y.M. Minsky and F.B. Schneider. Tolerating Malicious Gossip. *Distributed Computing*, 16:49–68, 2003.
- [26] S. Mitra. Iolus: A Framework for Scalable Secure Multicasting. *SIGCOMM Comput. Commun. Rev.*, 27(4):277–288, 1997.
- [27] Multicast Security. <http://datatracker.ietf.org/wg/msec/>
- [28] M. Onus and A.W. Richa. Minimum Maximum Degree Pub/Sub Overlay Network Design. In *INFOCOM 2009*, pages 882–890.
- [29] J. Pang and C. Zhang. How to Work with Honest but Curious Judges? In *Proc. 7th International Workshop on Security Issues in Concurrency*, pages 31–45, 2009.
- [30] S. Panjwani. Tackling Adaptive Corruptions in Multicast Encryption Protocols. In *TCC 2007*, pages 21–40.
- [31] A. Pelc. Fault-tolerant Broadcasting and Gossiping in Communication Networks. *Networks*, 28: 143–156, 1996.
- [32] A. Shamir. How to Share a Secret. *Communications of the ACM*, 22(11):612–613, 1979.
- [33] A.T. Sherman and D.A. McGrew. Key Establishment in Large Dynamic Groups Using One-Way Function Trees. *IEEE Transactions on Software Engineering*, 29(5):444–458, 2003.
- [34] D.R. Stinson. *Cryptography: Theory and Practice*, CRC Press, 3rd edition, 2005.
- [35] C.K.Wong, M. Gouda, and S. Lam. Secure Group Communications Using Key Graphs. *IEEE/ACM Transactions on Networking*, 8(1):16–30, 2000.
- [36] A.C. Yao. Protocols for Secure Computations. In *FOCS 1982*, pages 160–164.

APPENDIX A.
DETAILED PSEUDOCODE OF ALGORITHM CONGOS

Figure 1 depicts the interaction of the various services at a process i for a partition ℓ . Figures 2–5, present the detailed pseudocode of these services.

APPENDIX B.
FULL PROOFS

Proof of Theorem 1.: We may assume that n is sufficiently large (in fact, $n \geq 8$ is sufficient). Let c be a constant and x be a parameter, to be specified later. Suppose that only rumors with uniform deadlines $dmax$ are injected. Suppose that each process receives one rumor with a random set of destinations in the beginning of the computation, defined independently over processes and in such a way that for each process it is decided independently with probability x/n whether it belongs to this destination set (or not, otherwise). We argue that under this scenario (of the adversary), with probability at least $1 - x^{2c+2}/n^{c-1}$, no message can carry more than c rumors. Having this, observe that the number of pairs ($source_process, destination_process$) is at least $nx/2$, with probability at least $1 - e^{-nx/8} \geq 1 - 1/e$, by the Chernoff bound. It follows that the total number of rumor copies carried out by messages is at least $nx/2$. Therefore, the total number of messages needed to deliver all these rumors is at least $\frac{nx}{2c} = \Omega(nx)$, with probability at least $1 - 1/e - x^{2c+2}/n^{c-1}$.

Setting x to $n^{1/2-2/c}$, the above probability is at least $1 - 1/e - n^{-2} \geq 1 - 2/e > 0$. For any given ε , we can set $c = \lceil 2/\varepsilon \rceil$, which implies that the total number of messages needed to deliver the rumors is $\Omega(nx) \supseteq \Omega(n^{(3/2)-\varepsilon})$, with a positive probability. It follows, by applying the probabilistic argument, that there is a configuration of destination sets of size x such that $\Omega(n^{(3/2)-\varepsilon})$ messages is necessary, and this holds in any determined execution of any algorithm under such scenario. Hence, even randomized algorithms require total message complexity of $\Omega(n^{(3/2)-\varepsilon})$ with probability 1, which in view of the deadline $dmax$ gives $\Omega(n^{(3/2)-\varepsilon}/dmax)$ per-round message complexity with probability 1. Note that the suitable destination sets can be computed by the adversary up-front of the computation, the injection round is the same for all the rumors, and no crashes/restarts are needed to enforce this lower bound — hence the adversary is oblivious.

It remains to prove that with a positive probability, no message can carry more than c rumors, for some constant c , in the scenario of random destination sets, and to define parameters x and c . Consider any $c + 1$ processes. The probability that there are at least two processes in the intersection of destination sets of their rumors is at most

$$n^2 \cdot (x/n)^{c+1} \cdot (x/n)^{c+1} = x^{2c+2}/n^{2c}.$$

Here we used the union bound, where n^2 is the upper bound on the number of pairs in the destination set for the given set of $c + 1$ rumors and $(x/n)^{c+1} \cdot (x/n)^{c+1}$ is the probability that for a given such pair of different processes they are both in all destination sets. Therefore, with this probability, the rumors of these processes cannot be sent together in one message (as there is no other process in the common destination for these rumors). The number of such $(c + 1)$ -tuples of processes is at most n^{c+1} . Consequently, by the union bound, we get that with probability of at most

$$x^{2c+2}/n^{2c} \cdot n^{c+1} = x^{2c+2}/n^{c-1}$$

there are $c + 1$ processes with at least two processes in the intersection of the destination sets of their rumors. Recall that this event is a superset of the event that more than $c + 1$ rumors can be carried out by some message in the execution. \square

Proof of Theorem 2.: The proof of the theorem follows directly from the following two lemmas.

The first lemma states that confidentiality is not violated.

Lemma B1 (Confidentiality) *In any execution of algorithm CONGOS, and for any rumor ρ , if $q \notin \rho.D$, then at no point during the execution q learns $\rho.z$.*

Proof: Consider a rumor ρ injected at a processes p at round t of an execution of algorithm CONGOS. Also consider a process $q \notin \rho.D$. Assume that both p and q are continuously alive for the lifetime of rumor ρ (other cases, for example q not being continuously alive, are dealt in a similar fashion). We will show by investigating the flow of the algorithm that at no point will process q learn $\rho.z$, or be able to re-construct it.

Once ρ is injected in process p , based on the description of ConfidentialGossip at process p , and focusing on a partition ℓ , the rumor is split into $\rho_{0,\ell}$ and $\rho_{1,\ell}$. We consider the case where both processes are in the same group in partition ℓ . The other case is symmetric.

Without loss of generality, say that they belong in group $P_{0,\ell}$. Per the description of ConfidentialGossip, $\rho_{0,\ell}$ is disseminated, using GroupGossip $[\ell]$, only to processes in group $P_{0,\ell}$ (this is guaranteed by $Filter[\ell]$). Furthermore, since

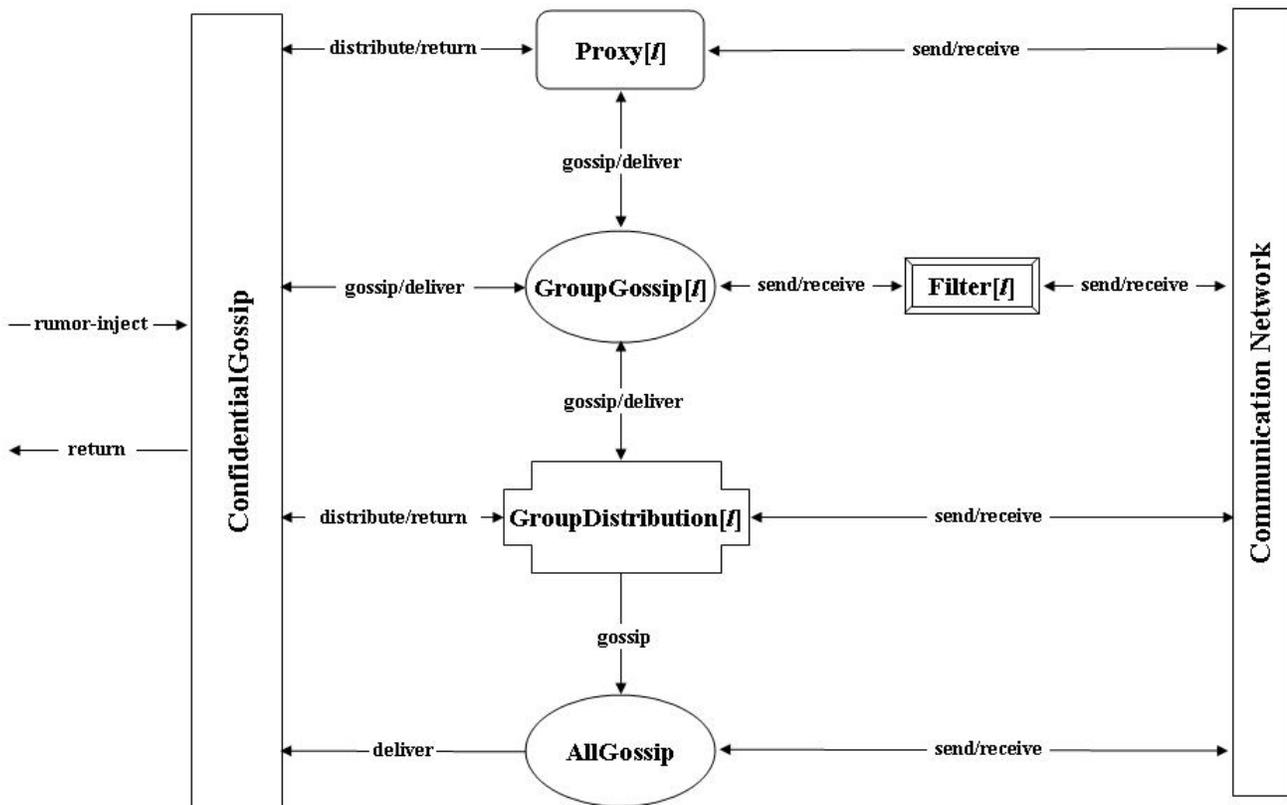


Figure 1. The interaction between the services at process i , for a partition ℓ .

service ConfidentialGossip($dline$) _{i}

```

1  state
2     $delivered-rumors, rumors-parts \subseteq R$  //  $R$  denotes the set rumors
3     $r, r0, r1 \in R$ 
4     $rumor-cache \subset R \times \mathbb{Z} \times \mathbb{Z}$ 
5     $hitSetM[1 \dots \log n][0 \dots 1]$  is a two dimensional matrix where each element is a tuple in  $[n] \times \mathbb{Z}$ .
6     $counter \in \mathbb{Z}$ 
7
8    //Time is divided into blocks of  $dline/4$  rounds.
9    //When a rumor is injected, in the first block it is split, in the second block the fragments are distributed via GroupGossip and the Proxy,
10   //in the third block the fragments are reassembled via GroupDistribution, and in the fourth block the sender receives confirmation.
11   //Upon a recovery from failure, the process retrieves the round number from the global clock and proceeds analogously.
12
13   input rumor-inject( $r$ ) //This marks the beginning of round 1 of a new block
14      $counter \leftarrow counter + 1$ 
15     for every  $\ell \in \{1, \dots, \log n\}$  do in parallel
16        $\langle r0, r1 \rangle \leftarrow \text{random-split}(r, counter, \sqrt{dline}, [n])$ 
17        $rumor-cache \leftarrow rumor-cache \cup \langle r, counter, round \rangle$ 
18       if  $i[\ell] = 0$  then //  $i[\ell]$  represents the  $\ell$ -th bit of the binary representation of  $i$ .
19         GroupGossip[ $\ell$ ].gossip( $r0$ ) //Gossip the rumor fragment in group
20         Proxy.distribute[ $\ell$ ]( $r1$ ) //Find a proxy to distribute the other rumor fragment in the other group
21       else
22         GroupGossip[ $\ell$ ].gossip( $r1$ )
23         Proxy[ $\ell$ ].distribute( $r0$ )
24
25   input Proxy[ $\ell$ ].return( $R$ )
26     GroupDistribution[ $\ell$ ].distribute( $R$ )
27
28   input GroupGossip[ $\ell$ ].deliver( $R$ )
29     GroupDistribution[ $\ell$ ].distribute( $R$ )
30
31   input GroupDistribution[ $\ell$ ].return( $R$ )
32      $rumor-parts \leftarrow rumor-parts \cup R$ 
33     for every  $r_1, r_2 \in rumor-parts$  do
34       if  $\text{merge}(r_1, r_2) = \langle \text{success}, r \rangle$  then
35         if  $r \notin delivered-rumors$  then
36            $delivered-rumors \leftarrow delivered-rumors \cup r$ 
37           return( $r$ )
38          $rumor-parts \leftarrow rumor-parts \setminus \{r_1, r_2\}$ 
39
40   input AllGossip.deliver( $R$ )
41     for every  $(\langle \text{distribution}, j, partition, h \rangle, *, *) \in R$  do
42        $hitSetM[partition, j[partition]] \leftarrow hitSetM[partition, j[partition]] \cup \{h\}$ 
43     for every  $\langle r, c, t \rangle \in rumor-cache$  do
44       if  $\exists \ell \in [1, \dots, \log n]$  where:
45          $\{(p_k, c) : p_k \in r.D\} \subseteq hitSetM[\ell, 0]$ 
46         and
47          $\{(p_k, c) : p_k \in r.D\} \subseteq hitSetM[\ell, 1]$ 
48       then  $rumor-cache \leftarrow rumor-cache \setminus \{\langle r, c, t \rangle\}$ 
49
50   In every round:
51     if  $\exists \langle r, c, t \rangle \in rumor-cache$  where  $round = t + r.d$  then
52       for every  $j \in r.D$  do
53         Network.send( $\langle \text{shoot}, r \rangle, i, j$ )
54
55   input Network.receive( $m, src, dest$ )
56     If  $m = \langle \text{shoot}, r \rangle$  then return( $r$ )

```

Figure 2. Main protocol at process i .

```

service Proxy( $dline, \ell$ ) $i$ 
1  state
2     $failed-proxies, current-proxies, proxy-ack, collaborators \subseteq [n]$ 
3     $status \in \{idle, active\}$ 
4     $my-rumors, waiting-rumors, proxy-buffer \subseteq R$ 
5     $r \in R$ 
6     $wakeup \in \mathbb{Z}$ 
7    //Time is divided into blocks of  $dline/4$  rounds.
8    //Each block is divided into iterations of  $\sqrt{dline} + 2$  rounds.
9    //Each iteration consists of 1 sending round, 1 gossip instance of  $\sqrt{dline}$  rounds, and 1 acknowledging round.
10   On recovery from failure:
11      $wake-up \leftarrow round$  // the value of round is retrieved from the global clock
12      $status \leftarrow idle$ 
13
14   At the beginning of round 1 of a new block:
15     if  $|round - wakeup| \geq dline/4$  then
16        $my-rumors \leftarrow waiting-rumors$ 
17        $waiting-rumors \leftarrow \emptyset$ 
18       if  $my-rumors \neq \emptyset$  then
19          $status \leftarrow active$ 
20          $failed-proxies, partial-rumors, proxy-buffer, proxy-ack \leftarrow \emptyset$ 
21          $collaborators \leftarrow \{j \in [n] : j[\ell] = i[\ell]\}$ 
22
23   At the beginning of round 1 of an iteration:
24     if  $status = active$  then
25        $current-proxies \leftarrow \Theta(n^{1+48/\sqrt{dline}} \log n) / |collaborators|$  processes chosen uniformly at random
26         from  $\{j \in [n] : j[\ell] \neq i[\ell]\} \setminus failed-proxies$ 
27       for every  $j \in current-proxies$  do Network.send( $\langle proxy, my-rumors \rangle, i, j$ )
28
29   At the beginning of round 2 of an iteration:
30      $collaborators \leftarrow \emptyset$ 
31     if  $status = active$  then GroupGossip[ $\ell$ ].gossip( $\langle proxy-buffer, failed-proxies, i \rangle, \sqrt{dline}, [n]$ )
32
33   At the beginning of the last round of an iteration:
34     if  $status = active$  then for every  $j \in proxy-ack$  do Network.send( $proxy-ack, i, j$ )
35
36   At the end of the last round of a block:
37     ConfidentialGossip[ $\ell$ ].return( $partial-rumors$ )
38
39   input ConfidentialGossip[ $\ell$ ].distribute( $r$ )
40      $waiting-rumors \leftarrow waiting-rumors \cup \{r\}$ 
41
42   input GroupGossip[ $\ell$ ].deliver( $R$ )
43     for every  $(\langle m, F, j \rangle, *, *) \in R$  do
44        $failed-proxies \leftarrow failed-proxies \cup F$ 
45       if  $status \neq idle$  then  $collaborators \leftarrow collaborators \cup \{j\}$ 
46        $partial-rumors \leftarrow partial-rumors \cup m$ 
47
48   input Network.receive( $\langle proxy, m \rangle, src, dest$ )
49      $proxy-buffer \leftarrow proxy-buffer \cup \{m\}$ 
50      $proxy-ack \leftarrow proxy-ack \cup \{src\}$ 
51
52   input Network.receive( $proxy-ack, src, dest$ )
53      $status \leftarrow idle$ 

```

Figure 3. Proxy search at process i for partition ℓ .

```

service GroupDistribution( $dline, \ell$ ) $i$ 
1  state
2     $partials, waiting-partials \subseteq [n] \times R$ 
3     $target-procs, collaborators, hitProcs \subseteq [n]$ 
4     $hitSet \subseteq [n] \times \mathbb{Z}$ 
5     $target-msg \subseteq R \times R \times \dots \times R \times [n]$ 
6     $wakeup \in \mathbb{Z}$ 
7     $status \in \{idle, active\}$ 
8    //Time is divided into blocks of  $dline/4$  rounds.
9    //Each block is divided into iterations of  $\sqrt{dline} + 2$  rounds.
10   //Each iteration consists of one initialization round, one distribution round and one gossip instance of  $\sqrt{dline}$  rounds.
11   On recovery from failure:
12      $wakeup \leftarrow round$ 
13      $status \leftarrow idle$ 
14
15   At the beginning of round 2 of a block:
16     if  $|round - wakeup| \geq 2dline/3$  then
17        $status \leftarrow active$ 
18        $partials \leftarrow waiting-partials$ 
19        $hitSet, waiting-partials \leftarrow \emptyset$ 
20        $collaborators \leftarrow \{j \in [n] : j[\ell] = i[\ell]\}$ 
21
22   At the beginning of round 2 of an iteration:
23     if  $status = active$  then
24        $hitProcs = \{p \in [n] : \langle p, \cdot \rangle \in hitSet\}$ 
25        $target-procs \leftarrow \Theta(n^{1+48/\sqrt{dline}} \log n / |collaborators|)$  processes chosen uniformly at random
26                                     from  $\{j \in [n] : j[\ell] \neq i[\ell]\} \setminus hitProcs$ 
27
28       for every  $j \in target-procs$  do
29          $target-msg \leftarrow \{r_k \in partials : j \in r_k.z.D\}$ 
30          $hitSet \leftarrow hitSet \cup \{\langle j, r_k.z.cnt \rangle : r_k \in target-msg\}$ 
31         Network.send( $\langle partials, target-msg \rangle, i, j$ )
32
33   At the beginning of round 3 of an iteration:
34      $collaborators \leftarrow \emptyset$ 
35     if  $status = active$  then GroupGossip[ $\ell$ ].gossip( $\langle share, hitSet, i \rangle, \sqrt{dline}, [n]$ )
36
37   At the end of the last round of a block:
38     AllGossip.gossip( $\langle distribution, i, \ell, hitSet \rangle, dline/4 - 1, [n]$ )
39
40   input ConfidentialGossip[ $\ell$ ].distribute( $r$ )
41      $waiting-partials \leftarrow waiting-partials \cup \{r\}$ 
42
43   input GroupGossip[ $\ell$ ].deliver( $\langle share, h, j \rangle$ )
44     if  $status = active$  then
45        $collaborators \leftarrow collaborators \cup \{j\}$ 
46        $hitSet \leftarrow hitSet \cup h$ 
47
48   input Network.receive( $m, src, dest$ )
49     if  $m = \langle partials, r_1, r_2, \dots \rangle$  then for every  $r_k \in m$  do ConfidentialGossip[ $\ell$ ].return( $r_k$ )

```

Figure 4. Rumor distribution between groups at process i for partition ℓ .

```

service Filter( $\ell$ ) $i$ 
1  input GroupGossip[ $\ell$ ].send( $m, src, dest$ )
2    if  $i[\ell] = src[\ell]$  then Network.send( $m, src, dest$ )
3
4  input Network.receive( $m, src, dest$ )
5    GroupGossip[ $\ell$ ].receive( $m, src, dest$ )

```

Figure 5. Filter ℓ at process i .

both p and q are continuously alive, GroupGossip $[\ell]$ guarantees (as shown in [12]) that q will indeed receive $\rho_{0,\ell}$. Similarly, $\rho_{1,\ell}$ is disseminated, using Proxy $[\ell]$, *only* to processes in group $P_{1,\ell}$. This can be seen by the description of Proxy $[\ell]$: For any iteration and any block, in the first round the processes in group $P_{0,\ell}$ send rumor fragments $x_{1,\ell}$ (including $\rho_{1,\ell}$) to processes in the other group (i.e., $P_{0,\ell}$) and then GroupGossip $[\ell]$ is used to disseminate these rumor fractions in group $P_{1,\ell}$ only. Hence, up to this point, process q knows only $\rho_{0,\ell}$.

Next, using the GroupDistribution $[\ell]$ service, processes in $P_{0,\ell}$ collaborate in sending (among other rumor fragments) $\rho_{0,\ell}$ to each process $w \in P_{1,\ell}$ where $w \in \rho.D$ (this knowledge is retrieved using the mentioned metadata). Similarly, the processes in $P_{1,\ell}$ disseminate $\rho_{1,\ell}$ to processes in $P_{0,\ell}$ that belong in $\rho.D$. By assumption $q \notin \rho.D$ and hence no process in $P_{1,\ell}$ ever sends $\rho_{1,\ell}$ to process q . The exchange of *hitSets* does not reveal any information on the rumor fragments and so does not the instance of AllGossip used in the last round of the last iteration of each block in GroupDistribution $[\ell]$. So process q still knows only $\rho_{0,\ell}$.

Finally, if the deadline of rumor ρ is about to expire, and there is no confirmation that ρ has been delivered (see the last two bullets in the outline of ConfidentialGossip), then the process p sends ρ directly to every process in $\rho.D$. Again, since $q \notin \rho.D$, p does not send ρ to q . Note that after this phase or if p receives a confirmation that ρ has been delivered, no further message is exchanged with respect to ρ and hence q never gets to learn ρ . As q cannot construct ρ only from $\rho_{0,\ell}$ or any combination of different partitions of ρ in the various group partitions (that run in parallel), the thesis of the lemma follows. \square

We now show that algorithm CONGOS delivers admissible rumors before they expire.

Lemma B2 (Correctness wrt QoD) *In any execution of algorithm CONGOS, and for any rumor ρ injected at a process p at round t of the execution, if p and $q \in \rho.D$ are continuously alive for the lifetime of the rumor, then q learns ρ by round $t + \rho.D$.*

Proof: Fix $\rho = \langle z, d, D \rangle$ injected at process p at round t , and consider process $q \in \rho.D$. Both p and q are continuously alive for the rumors lifetime. Hence ρ is admissible for q . We show that q will learn, with probability 1, rumor ρ by time $t + \rho.d$.

From the last two bullets in the outline of ConfidentialGossip we have that if the deadline of rumor ρ is about to expire, and there is no confirmation that ρ has been delivered, then process p sends ρ directly to every process in $\rho.D$, including q . Hence, what remains to be proved is that if p receives confirmation that ρ was delivered before the rumor has expired, then q has indeed learned ρ (i.e., it has learned both fragments of ρ in some partition ℓ). Process p will confirm that ρ has been delivered, if a message from AllGossip confirms that, for some partition ℓ , both fragments of a rumor ρ have been sent to every destination in $\rho.D$ (including q). Then the thesis of the lemma follows from these observations:

- No process will include q in its *hitSet* if it has not sent its partial rumor to q in the GroupDistribution service of some partition ℓ .
- Say w is a process that has included q in its *hitSet*. The dissemination of *hitSet* takes place after rumor fragments to the processes being “hit” are sent. If w would fail prior than hitting q , then its *hitSet* containing q would not have been disseminated through the AllGossip service. Hence, when p receives a *hitSet* (from AllGossip) containing q , it is the case that q was indeed hit.
- Messages sent from non-faulty processes to non-faulty processes are not lost (unless they go through the filter, which is not the case here).

This completes the proof. \square

Proof of Lemma 5.: We first observe that, throughout a block, every process p_i that is sending messages remains in the *collaborators* set for every other p_j in the same group of partition ℓ : Initially, *collaborators* contains every process; in each iteration. If p_i and p_j both remain alive to proceed in a later iteration, it means they both sent a rumor in the previous iteration indicating that they were alive. By the guarantees of GroupGossip, this rumor must have been delivered, and hence $p_i \in \text{collaborators}_j$ (and *vice versa*). Notice that if a process is restarted, it does not begin sending messages again until the next block.

In Proxy $[\ell]$ and GroupDistribution $[\ell]$, each process sends $n^{1+48/\sqrt{dline}} \log n / |\text{collaborators}|$ messages in, respectively, the first and second round of an iteration. The bound on *collaborators* implies the desired result.

In Proxy $[\ell]$, each process that received a proxy request sends a response at the end of an iteration. Each response is the result of an earlier request in the first round of the iteration, and we have already bounded the message complexity of the first round of an iteration, leading here too to a bound of $O(n^{1+48/\sqrt{dline}} \log n)$. \square

Proof of Lemma 6.: Fix ℓ to be the partition identified in Lemma 3 such that at least one process remains alive throughout the interval in both groups of the partition. Assume, without loss of generality, that $p_i \in P_{0,\ell}$. (The alternate case is symmetric.) Since p_i injects rumor fragment $\rho_{0,\ell}$ into the GroupGossip[ℓ] service with deadline \sqrt{dline} , it is guaranteed to reach every process in $P_{0,\ell}$ that remains alive throughout the interval. It remains to analyze the behavior of the Proxy[ℓ] service, focusing on the first complete block after rumor ρ is injected beginning at time $t + dline/4 - (t \bmod dline)$.

We need to show that each process in $P_{0,\ell}$ succeeds in finding a proxy in $P_{1,\ell}$. Let $Z = n^{48/\sqrt{dline}}$. We will argue that in every pair of iterations, one of the three following events occurs:

- 1) At least a $(1 - 1/Z)$ fraction of processes in $P_{0,\ell}$ that were alive at the beginning of the first iteration fail by the end of the second iteration.
- 2) At least a $(1 - 1/Z)$ fraction of processes in $P_{1,\ell}$ that were alive at the beginning of the first iteration fail by the end of the second iteration.
- 3) In the second iteration, at least a $(1 - 1/Z)$ fraction of processes in $P_{0,\ell}$ succeed in finding a proxy.

Notice that from this claim, we can conclude that by the end of $3 \log_Z(n)$ pairs of iterations, either every process in one of the two groups has failed, or every process in $P_{0,\ell}$ has succeeded in finding a proxy. Since by assumption there is at least one process alive in both groups of partition ℓ , and since $\log_Z(n) = \sqrt{dline}/48$, we conclude that by the end of $6 \log_Z n \leq \sqrt{dline}/8$ iterations, every process in $P_{0,\ell}$ has succeeded in finding a proxy. And once a process has succeeded in finding a proxy, it follows from the guarantees of GroupGossip that its rumor fragment is distributed to every non-failed process in the other group.

We now argue that in each iteration, one of the three events described above occurs. Fix a pair of iterations, let A be the set of processes in $P_{0,\ell}$ still active at the beginning of the first iteration, and let B be the set of processes in $P_{1,\ell}$ still active at the beginning of the first iteration. Assume the neither event (1) nor event (2) occur, i.e., that at the end of the iteration there are at least $|A|/Z$ processes still active in $P_{0,\ell}$ and $|B|/Z$ processes still active in $P_{1,\ell}$.

We now calculate the probability that a process in A successfully finds a proxy in B that does not fail by the end of the second iteration. In the second iteration, each process in A has an estimate of the size of A that is at most $|A|$. Thus, every process in A sends at least $c(n/|A|) \cdot Z \log n$ proxy requests, for some constant c . Moreover, in the first iteration, each process adds every process not in B to its set of *failed-proxies*, so we can conclude that every proxy request is sent to some process in B .

Fix some subset $A' \subseteq A$ of size at least $|A|/Z$, and fix some subset $B' \subseteq B$ of size at least $|B|/Z$. We now calculate the probability that a given process in A' fails to send a message to some process in B' as:

$$\begin{aligned} (1 - |B'|/|B|)^{cnZ \log n/|A|} &\leq (1 - 1/Z)^{cnZ \log n/|A|} \\ &\leq (1/e)^{cn \log n/|A|} \end{aligned}$$

The probability that *every* one of the at least $|A|/Z$ processes in A' fails to send a message to some process in B' is at most $e^{-cn \log n/Z}$.

We now take a union bound over all possible sets A' and B' . Specifically, there are at most $|A|^{|A|/Z} \leq e^{(|A|/Z) \log |A|}$ possible sets A' ; there are at most $|B|^{|B|/Z} \leq e^{(|B|/Z) \log |B|}$ possible sets B' . Thus, the probability that there exists any set A' and any set B' where every process in A' misses every process in B' is at most:

$$\begin{aligned} \frac{e^{(|A|/Z) \log |A| + (|B|/Z) \log |B|}}{e^{c(n/Z) \log n}} &\leq \frac{e^{(n/Z) \log n + (n/Z) \log n}}{e^{c(n/Z) \log n}} \\ &\leq e^{-(c-2)(n/Z) \log n} \end{aligned}$$

Thus, for sufficiently large c , with high probability at most $|A|/Z$ processes do not successfully find proxies, as required. \square

Proof of Lemma 7.: By Lemma 6, we know that if there is at least one process alive in both groups of partition ℓ , then every active process has received an appropriate rumor fragment. It remains to show that during the following block of rounds, for every process $p_j \in \rho.D$, at least one process from each group sends its rumor fragment to p_j . Fix some process $p_j \in \rho.D$ for the remainder of the proof. Without loss of generality, we focus on process in $P_{0,\ell}$; the case for the other group is symmetric. We now examine the behavior of the GroupDistribution service.

Let $Z = n^{48/\sqrt{dline}}$. We argue that in each pair of iterations of the GroupDistribution[ℓ] service, one of the following two events occurs:

- 1) At least a $(1 - 1/Z)$ fraction of processes in $P_{0,\ell}$ that were alive at the beginning of the first iteration fail by the end of the second iteration.

- 2) For every process p_k active throughout both iterations, the set of processes $[n] \setminus \text{hitProcs}_k$ decreases by a factor of Z by the end of the second iteration. ($\text{hitProcs}_k = \{p_q \in [n] : \langle p_q, \cdot \rangle \in \text{hitSet}_k\}$.)

From this claim, we conclude that within $2 \log_Z n$ pairs of iterations, either every process in $P_{0,\ell}$ fails, or every process has been added to hitProcs . Since a process is only added to hitProcs after it has been sent all the available rumor fragments in $P_{0,\ell}$, and since we have assumed that at least one process in $P_{0,\ell}$ remains alive throughout the block, we conclude that by the end of $4 \log_Z n \leq \sqrt{\text{dline}}/8$ iterations, every process has been sent all of its rumor fragments.

We now proceed to prove that one of the two above events occurs. Fix a pair of iterations, let A be the set of processes in $P_{0,\ell}$ still active at the beginning of the first iteration, and let H be the set of processes in $[n] \setminus \text{hitProcs}_k$ at the beginning of the first iteration, for the process p_k with the largest set hitProcs_k at the beginning of the first iteration, where p_k is active through both iterations. Assume that event (1) does not occur, i.e., there are at least $|A|/Z$ processes still active in $P_{0,\ell}$ at the end of the second iteration.

In the second iteration, each process in A has an estimate of the size of A that is at most $|A|$. Thus, every process in A sends at least $c(n/|A|) \cdot Z \log n$ messages. Moreover, during the first iteration, processes share their hitSets , and hence in the second iteration, messages are only sent to processes that were not in hitSet_k .

For a given subset $A' \subseteq A$ of $|A|/Z$ processes, for a given subset $H' \subseteq H$ of $|H|/Z$ processes, we calculate the probability that no process in A' sends a message to some process in H' :

$$(1 - 1/Z)^{c(n/|A|)(|A|/Z) \cdot Z \log n} \leq (1/e)^{c(n/Z) \log n}$$

There are at most $\binom{|A|}{|A|/Z} \leq e^{(|A|/Z) \log |A|}$ possible subsets A' , and at most $\binom{|H|}{|H|/Z} \leq e^{(|H|/Z) \log |H|}$ subsets H' . Thus, by a union bound over all possible subsets, the probability that *any* subset of $|A|/Z$ processes does not hit some subset of $|H|/Z$ processes is at most:

$$\frac{e^{(|A|/Z) \log |A| + (|H|/Z) \log |H|}}{e^{c(n/|A|) \log n}} \leq \frac{e^{2(n/Z) \log n}}{e^{c(n/Z) \log n}} \leq e^{-(c-2)(n/Z) \log n}$$

From this we conclude that the set H decreases by a factor of Z with high probability, concluding our proof. \square

APPENDIX C. SERVICE OUTLINES FOR ALGORITHM WITH COLLUSIONS

Differences from algorithm CONGOS are included in a box and annotated with **boldface text**.

Outline of ConfidentialGossip service at p_i :

- Do in parallel for each $\ell = 1, \dots, c\tau \log n$:
 - 1) Split rumor ρ into a **sequence** $\langle \rho_{0,\ell}, \rho_{1,\ell}, \dots, \rho_{\tau,\ell} \rangle$.
 - 2) If p_i is in group b of partition ℓ , inject $\rho_{b,\ell}$ into $\text{GroupGossip}[\ell]$, and inject **all** $\rho_{a,\ell}, a \neq b$, into $\text{Proxy}[\ell]$. Together, these two services ensure that each rumor fragment is delivered to every non-failed process in the appropriate group of the partition.
 - 3) For each rumor fragment received from $\text{GroupGossip}[\ell]$ or $\text{Proxy}[\ell]$, inject the fragment into $\text{GroupDistribution}[\ell]$.
 - 4) Save every fragment received from $\text{GroupDistribution}[\ell]$, and reassemble and deliver rumors as fragments become available.
 - Whenever a message from AllGossip confirms that, for some partition ℓ , **all** $\tau + 1$ fragments of a rumor ρ , initiated at p_i , have been sent to every destination in $\rho.D$, confirm that ρ has been delivered.
 - Whenever a deadline is about to expire for some rumor ρ initiated at p_i , and there is no confirmation that ρ has been delivered, send ρ directly to every process in $\rho.D$.
-

Outline of Proxy[ℓ] at p_i :

- Time is divided into blocks of length $dline/4$.
 - At the beginning of a block, collect all the fragments that have been injected since the last block began, and set status to active.
 - Each block is divided into *iterations* of $\sqrt{dline} + 2$ rounds. In each iteration, we maintain a set *collaborators* of the active processes in the same group as p_i . We also keep track of *failed-proxies*, i.e., those that we have already learned (in previous iterations) have failed in this block. For each iteration, repeat:
 - Round 1: **for each other group**, send every rumor fragment associated with that group to $n^{1+48/\sqrt{dline}} \log n / |collaborators|$ processes chosen uniformly at random from that group, excluding processes in *failed-proxies*. (Notice that as long as the set *collaborators* is a good estimate of the set of collaborators, this ensures a good bound on the message complexity of this step.) Every process that receives a request to be a proxy for **some other group(s)** caches the received rumor fragments.
 - Rounds 2, \dots , $\sqrt{dline} + 1$: initiate a GroupGossip[ℓ] in which processes in the same group as p_i share the set of *failed-proxies*, as well as establish the set of *collaborators*, i.e., members of the group that are still active. Processes also share all the rumor fragments received from the other **groups**. (The deadline for rumors in GroupGossip[ℓ] here is \sqrt{dline} .)
 - Round $\sqrt{dline} + 2$: Any process that was asked to be a proxy for the other **groups** sends an acknowledgment that proxying was successful. Any process that sent a request, and does not receive an acknowledgment, adds the non-acknowledging processes to the set of *failed-proxies*.
-

Outline of GroupDistribution[ℓ] at p_i :

- Time is divided into blocks of length $dline/4$.
 - At the beginning of the second round of a block, collect all the fragments that have been injected since the first round of the block, and set status to active. (The first round of the block is spent waiting for rumor fragments from the previous block.)
 - Each block is divided into *iterations* of $\sqrt{dline} + 2$ rounds. In each iteration, we maintain a set *collaborators* of the active processes in the same group as p_i . We also keep track of a set *hitSet* of processes that have been sent a message in this block. Each process in this set was sent all the rumor fragments for this block. For each iteration, repeat:
 - Round 1: wait for rumor fragments to be injected.
 - Round 2: **for each other group**, send every “appropriate” rumor fragment to $n^{1+48/\sqrt{dline}} \log n / |collaborators|$ processes chosen uniformly at random from that other group, excluding processes in *hitSet*. By appropriate we mean that if p_j is a process chosen randomly by p_i , then p_i sends to p_j only the rumor fragments in which p_j is in the destination set. (Recall that each partial rumor contains the target destination set as part of the metadata.) Every process that receives rumor fragments can now reconstruct the rumor and return it to its user (via the ConfidentialGossip service).
 - Rounds 3, \dots , $\sqrt{dline} + 2$ rounds: initiate an instance of GroupGossip[ℓ] (with deadline \sqrt{dline}) in which processes in the same group as p_i share their *hitSets*, as well as count how many members of the group are still active.
 - In the last round of the block, initiate an instance of AllGossip (with deadline $dline/4 - 1$). Each process p_i gossips the information in its *hitSet*, but without including the rumor fragments themselves. That is, if the *hitSet* of process p_i indicates that some rumor fragment $\rho_{0,\ell}$ was sent to some process p_j , and if $\rho_{0,\ell}$ has identifier r , then p_i gossips that the fragment 0 for partition ℓ of the rumor associated with identifier r was sent to p_j . This provides sufficient information for the source to determine whether the rumor was delivered, without revealing the contents of the rumor. (See the description of the ConfidentialGossip service, above, for how this information is used.)
-