Randomized Rumor Spreading

R. Karp* C. Schindelhauer† S. Shenker‡ B. Vöcking§

Abstract

We investigate the class of so-called epidemic algorithms that are commonly used for the lazy transmission of updates to distributed copies of a database. These algorithms use a simple randomized communication mechanism to ensure robustness. Suppose n players communicate in parallel rounds in each of which every player calls a randomly selected communication partner. In every round, players can generate rumors (updates) that are to be distributed among all players. Whenever communication is established between two players, each one must decide which of the rumors to transmit. The major problem (arising due to the randomization) is that players might not know which rumors their partners have already received. For example, a standard algorithm forwarding each rumor from the calling to the called players for $\Theta(\ln n)$ rounds needs to transmit the rumor $\Theta(n \ln n)$ times in order to ensure that every player finally receives the rumor with high probability.

We investigate whether such a large communication overhead is inherent to epidemic algorithms. On the positive side, we show that the communication overhead can be reduced significantly. We give an algorithm using only $O(n \ln \ln n)$ transmissions and $O(\ln n)$ rounds. In addition, we prove the robustness of this algorithm, e.g., against adversarial failures. On the negative side, we show that any address-oblivious algorithm (i.e., an algorithm that does not use the addresses of communication partners) needs to send $\Omega(n \ln \ln n)$ messages for each rumor regardless of the number of rounds. Furthermore, we give a general lower bound showing that time- and communication-optimality cannot be achieved simultaneously using random phone calls, that is, every algorithm that distributes a rumor in $O(\ln n)$ rounds needs $\omega(n)$ transmissions.

1 Introduction

We investigate the problem of spreading rumors in a distributed environment using randomized communication. Suppose n players exchange information in parallel communication rounds over an indefinite time. In each round t, the players are connected by a communication graph G_t generated by *random phone calls* as follows: each player u selects a communication partner v at random and u *calls* v; two players u and v are connected by an edge in G_t if u calls v in round t. Rumors can be started in any round by any player and can be transmitted in both directions along the edges in the graph G_t in round t. The goal is to spread the rumor among all participating players using a small number of rounds and a small number of transmissions.

The motivation for using randomized communication is that it naturally provides robustness, simplicity, and scalability. For example, consider the following so-called *push algorithm*. Starting with the round in which a rumor is generated, each player that holds the rumor forwards it to a communication partner selected independently and uniformly at random. The distribution of the rumor is terminated after some fixed number of $O(\ln n)$ rounds. At this time all players are informed, with high probability[1].

Clearly, one can also inform all players in $O(\ln n)$ rounds using a deterministic interconnection of constant degree, e.g., a shuffle network. (For an overview of deterministic information dissemination we refer to [5] or [6].) The advantage of the randomized push algorithm, however, is its inherent robustness against several kinds of failures compared to deterministic schemes that either need substantially more time [4] or can tolerate only a relatively small number of faults [10]. For example, consider node failures in which a player (different from the player starting the rumor) fails to communicate or simply crashes and forgets its rumors. Obviously, when using a sparse deterministic network, even a single node failure can result in a large fraction of play-

*Email: karp@icsi.berkeley.edu. International Computer Science Institute, Berkeley and University of California at Berkeley.

†Email: schindel@tcs.mu-luebeck.de, Institut für Theoretische Informatik, Med. Universität zu Lübeck. Parts of this work are supported by a stipend of the “Gemeinsames Hochschulsonderprogramm III von Bund und Länder” through the DAAD.

‡Email: shenker@icsi.berkeley.edu. International Computer Science Institute, Berkley.

§Email: voecking@mpi-sb.mpg.de. Max-Planck-Institut für Informatik, Saarbrücken, Germany.

[1] The term with high probability (w.h.p.) means with probability at least $1 - O(n^{-\alpha})$ for an arbitrary constant $\alpha > 0$.

ers not receiving the rumor. When using the randomized push algorithm, however, the effects of node failures are very limited. In fact, it is not difficult to prove that F node failures (specified by an oblivious adversary) result in only $O(F)$ uninformed players, w.h.p.

Unfortunately, the push algorithm produces a large communication overhead. In fact, it needs to forward each individual rumor $\Theta(n \ln n)$ times before all players are informed, in comparison to a deterministic scheme which requires only $n - 1$ transmissions. It seems that the large number of transmissions is the price for the robustness. This gives rise to the question whether this additional communication effort is a special property of the above push algorithm or is inherent to rumor spreading using random phone calls in general.

1.1 Background

Demers et al. [2] introduced the idea of using so-called epidemic algorithms for the lazy update of data objects in a data base replicated at many sites, e.g., yellow pages, name servers, or server directories. In particular, they propose the following two concepts:

- *Anti-entropy:* Every site regularly chooses another site at random and resolves all differences by exchanging the complete data base contents.

- *Rumor mongering:* When a site receives a new update it becomes a "hot rumor". While a site holds a "hot rumor", it periodically chooses another site at random and sends the rumor to the other site.

It turns out that anti-entropy is extremely reliable but produces such an enormous amount of communication that it cannot be used too frequently. The idea of rumor mongering is to exchange only recent updates, thereby reducing the communication overhead significantly. In practice one might use a combination of both concepts, that is, using rumor mongering frequently and anti-entropy very rarely in order to ensure that all updates are recognized by all sites. In this paper, we solely investigate algorithms implementing the rumor mongering concept.

The original idea for rumor spreading was to send rumors only from the caller to the called player (*push transmission*) [2]. Several termination mechanisms deciding when a rumor becomes "cold" so that it transmission is stopped were investigated. All these algorithms share the same phenomenon: the fraction u of players that do not know a particular rumor decreases exponentially with the number of transmissions t (i.e., messages that contain this rumor). So-called *mean field equations* (implicitly assuming that u is sharply concentrated around its mean value $\mathbf{E}[u]$) lead to the conjecture that $u \approx \exp(-t/n)$ for all variants of the push algorithm that have been investigated. In other words, a push algorithm needs $\Theta(n \ln n)$ transmissions for sending a rumor to all players.

A further idea introduced in [2] is to send rumors from the called to the calling player (*pull transmission).* It was observed that the number of uninformed players decreases much faster using a pull scheme instead of a push scheme. This kind of transmission makes sense if updates occur frequently so that (almost) every player places a random call in each round anyway. Mean field equations lead to the conjecture that $u \approx \exp(-2^t)$ for pull schemes. Clearly, this double exponential behavior implies that only $\Theta(n \ln \ln n)$ transmissions are needed if the distribution of the rumor can be stopped at the right time. Such a termination mechanism, however, is not presented. Instead, the authors predict that $\Theta(n \sqrt[3]{\ln n})$ transmissions are sufficient for some other specific termination mechanisms.

The work of Demers et al. initiated an enormous amount of experimental and conceptual study of epidemic algorithms. For example, there is a variety of research issues like consistency, correctness, data structures, and efficiency [1, 7, 8, 9, 12]. Recent theoretical work concentrates on the robustness against Byzantine failures [11]. In this paper, we concentrate only on the efficiency of these randomized algorithms. In particular, we study their time and communication complexity using a simple model for the underlying randomized communication.

1.2 The random phone call model

Let V denote the set of players. The communication graph $G_t = (V, E_t \subseteq V \times V)$ of round $t \geq 1$ is obtained by a distributed, randomized process. In each round, each player u chooses a communication partner v from V at random and u *calls* v. Unless otherwise stated, we assume that all players choose their communication partners independently and uniformly at random from V.

Even though we envisage an application (such as the lazy transmission of updates to distributed copies of a database) in which rumors are constantly generated by different players, our analysis is concerned with the distribution of a single rumor only. We focus on the lifetime of the rumor and the number of transmissions rather than the number of connections established because the latter cost is amortized over all the rumors using that connection.

In round t, the rumor and other information can be exchanged in both directions along the edges of G_t. Whenever a connection is established between two players, each one of them (if holding the rumor) has to decide whether to transmit the rumor to the other player, typically without knowing whether this player has received the rumor already. Regarding the flow of information, we distinguish between push and pull transmissions. Assume player u calls player v.

- The rumor is *pushed* If u tells v the rumor.
- The rumor is *pulled* if v tells u the rumor.

We do not limit the size of the information exchanged in any way. Each information exchange between neighboring players in a round is counted as a single transmission. (We point out that our algorithms only add small counter values to rumors, whereas our lower bounds hold even for algorithms in which players exchange their complete history whenever the rumor is sent in either direction.) Communication inside each round, however, is assumed to proceed in parallel, that is, any information received in a round cannot be forwarded to another player in the same round.

The major issue that has to be specified by a rumor spreading algorithm is how players decide whether the rumor shall be forwarded to a communication partner. An algorithm is called *distributed* if players make these decisions using only local knowledge. In other words, the decision whether a player sends a message to a communication partner in round t depends only on the player's *state* in that round. The initial state of a player is defined by the player's address, the number of players, and possibly a random bit string. The state of a player in round $t \geq 1$ is a function of its initial state, the addresses of its neighbors in the communication graphs $G_1, \ldots, G_t$, and the information received in rounds 1 to $t-1$. (For our lower bounds we allow the state to depend in addition on a globally known round number and the birth date of the rumor considered.)

Finally, an algorithm is called *address-oblivious* if a player's state in round t does not depend on the addresses of the neighbors in G_t but only on the number of neighbors in G_t. (The state can still depend on the addresses of neighbors in $G_1, \ldots, G_{t-1}$.) We point out that all rumor spreading algorithms proposed by Demers et al. [2] are address-oblivious.

1.3 New results

We prove that the number of transmissions can be reduced significantly when the rumor is sent in both directions, that is, when using push and pull rather than only push operations. We introduce a *simple push&pull algorithm* spreading the rumor to all players in $O(\ln n)$ rounds using only $O(n \ln \ln n)$ transmissions in comparison to $\Theta(n \ln n)$ as used by the push algorithm.

The drawback of the simple push&pull-algorithm is that its success heavily relies on a very exact, global estimation of the right termination time. This mechanism is very sensitive to any kind of errors that influence the expansion of the set of informed players. In order to improve the robustness, we devise a distributed termination scheme, called the *median-counter algorithm*, that is provably robust against adversarial node failures as well as stochastic inaccuracies in establishing the random connections.

In particular, we show that the efficiency of the algorithm does not rely on the fact that players choose their communication partners uniformly from the set of all players. We show that the median-counter algorithm takes $O(\ln n)$ rounds and needs only $O(n \ln \ln n)$ transmissions regardless of the probability distribution used for establishing the random connections as long as all players act independently and each player uses the same distribution $\mathcal{D} : V \rightarrow [0, 1]$ to select its communication partner. For example, this allows sampling from an arbitrary address directory (possibly with redundant addresses and some non-listed players as in a telephone book). In other words, the algorithm can be executed even without global knowledge about the set of players.

In addition, we provide lower bounds assuming that the communication partners are selected using the uniform probability distribution. Both the simple push&pull algorithm as well as the median-counter algorithm are address-oblivious and use only $O(n \ln \ln n)$ transmissions. We prove a corresponding lower bound showing that any address-oblivious algorithm needs to perform $\Omega(n \ln \ln n)$ transmissions in order to inform all players. We point out that this bound holds independently of the number of rounds executed.

The situation changes substantially when considering general (i.e., possibly non-address-oblivious) algorithms. Allowing $\Theta(n \ln n)$ rounds, an algorithm that exploits the addresses of communication partners can spread the rumor using only $n-1$ transmissions. Here is a simple example. The player initiating the rumor simply waits until each other player appears as communication partner for the first time and then forwards the rumor to this player. Clearly, this is not a practical algorithm as it takes too many rounds. Nevertheless, it illustrates the additional possibilities when the addresses of communication partners can be exploited.

The above example leads to the question of whether general epidemic algorithms can spread a rumor in a small number of rounds while using only a linear number of transmissions. We give a lower bound answering this question negatively. In particular, we show that any randomized rumor spreading algorithm running for $O(\ln n)$ rounds requires $\omega(n)$ transmissions. This lower bound holds regardless of the amount of information that can be attached to the rumors. For example, players might always exchange their complete communication history whenever the rumor is transmitted in either direction. Thus, there is a fundamental gap between rumor spreading algorithms based on random interconnections and deterministic broadcasting schemes.

2 The advantage of push&pull

First, let us explain the differences in the propagation of the rumor obtained by push transmissions on the one hand and pull transmissions on the other hand.

- Consider a *push scheme* in which every informed player, in every round, forwards the rumor to the player it calls until all players are informed. In this case the set of informed players grows exponentially until about $n/2$ players are informed. At about this time the exponential growth of the set of informed players stops. Starting from this point of time, let us consider the set of uninformed players. Once half of the players are informed, this set shrinks by a constant factor in each round. At the end of the rumor spreading process this factor is about $1 - 1/\mathrm{e}$ since the fraction of players that do not receive a call in a round is about $1/\mathrm{e}$. Thus, the shrinking phase takes $\Theta(\ln n)$ rounds until every player has received the rumor, and the push algorithm sends $\Theta(n)$ messages in each of these rounds.

- Now consider a *pull scheme* in which only called players send the rumor towards the calling players. In this case, the player starting the rumor may have to wait some rounds until it is called for the first time so that the propagation in the first rounds becomes unpredictable. But eventually (after $O(\ln n)$ rounds, w.h.p.) about $n/2$ of the players will be informed. From this time on, the pull algorithm has an advantage against the push algorithm as the fraction of uninformed players roughly squares from round to round. This is because in a round starting with ϵn uninformed players, each individual player has probability $1 - \epsilon$ to receive the rumor, so that the probability of staying uninformed is ϵ, resulting in an expected number of $\epsilon^2 n$ uninformed players at the end of the round. Thus, we can expect that the shrinking phase only takes $\Theta(\ln \ln n)$ rounds so that only $\Theta(n \ln \ln n)$ messages are sent during this phase.

In order to combine the predictability of the push scheme with the quadratic-shrinking property of the pull scheme, we simply send the rumor in both directions whenever possible. In detail, our *push&pull scheme* works as follows. The creator of the rumor initiates a time-counter with 0 representing the *age* of the rumor. The age is incremented in every round and distributed with the rumor. In every round every informed player pushes and pulls unless the age of the rumor is higher than $t_{\max} = \log_3 n + O(\ln \ln n)$. In the following theorem, we assume the uniform distribution and a perfect interconnection without failures.

Theorem 2.1 *The simple push&pull-scheme informs all players in time* $\log_3 n + O(\ln \ln n)$ *using* $O(n \ln \ln n)$ *messages w.h.p.*

Proof. Let S_t be the set of informed players and U_t the set of uninformed players at the end of round t. Define $s_t = |S_t|$ and $u_t = |U_t|$. We distinguish four consecutive phases.

1. The *startup phase* starts in the round in which the rumor is created and ends with the first round after whose execution there are at least $(\ln n)^4$ informed players for the first time. At the beginning of the first round only one player holds the rumor. If we execute c rounds then the probability that this player has at least once called an uninformed player (i.e., did not call itself) is $1 - n^{-c}$. Thus, we double the number of informed players in c rounds, w.h.p. In general, starting with at most $(\ln n)^4$ informed players, we need at most c rounds to double the number of informed players, w.h.p. Thus $O(\ln \ln n)$ rounds are sufficient to achieve $(\ln n)^4$ informed players.

2. The *exponential-growth phase* ends with the round after whose execution there are at least $n/\ln n$ informed players for the first time. The expected number of messages (containing the rumor) sent during round t in this phase is $2s_{t-1}$ because each player holding the rumor calls one player and is called by one player on average. Applying a Chernoff bound yields that the number of messages actually sent is $m = (2 \pm o(1/\ln n))s_{t-1}$, w.h.p, applying $s_{t-1} \geq (\ln n)^4$. (Due to space limitations, we dot not explain the mathematical details behind the application of Chernoff bounds in this extended abstract.) Unfortunately, some of these messages are *wasted* as they are directed to the same player or an informed player. As interconnections are chosen at random, the probability that a particular message is wasted is at most $s_{t-1}/n + m/n$. This expression is bounded above by $(3 + o(1/\ln n))/\ln n$ because $s_{t-1} \leq n/\ln n$. As a consequence,

$$\begin{aligned} \mathbf{E}\left[s_t\right] &= s_{t-1} + m\left(1 - \frac{3 + o(1/\ln n)}{\ln n}\right) \\ &= s_{t-1}\left(3 - O(1/\ln n)\right) \ . \end{aligned}$$

   Applying a Chernoff bound yields

   $$s_t = (1 \pm o(1/\ln n))\mathbf{E}\left[s_t\right] = s_{t-1}\left(3 \pm O(1/\ln n)\right) \ ,$$

   since $\mathbf{E}\left[s_t\right] \geq (\ln n)^4$. Assuming this expansion factor in each round, we can observe that this phase takes $\log_3 n \pm O(\ln \ln n)$ rounds.

3. The *quadratic-shrinking phase* ends with the round after whose execution there are at least $\sqrt{n}(\ln n)^4$ uninformed players for the last time. Even if we only take into account pull transmissions we obtain (by following the arguments explaining the general properties of pull algorithms) that

   $$\mathbf{E}\left[\frac{u_t}{n}\right] \leq \left(\frac{u_{t-1}}{n}\right)^2 .$$

Applying a Chernoff bound yields

$$u_t \leq \left(1 + \frac{1}{\ln n}\right) \frac{(u_{t-1})^2}{n} \ ,$$

w.h.p., provided $u_t \geq \sqrt{n}(\ln n)^4$. Now some easy calculations show that we need $O(\ln\ln n)$ rounds until the number of uninformed players drops from $n/\ln n$ to $\sqrt{n}(\ln n)^4$.

4. In the *final phase*, we inform the few remaining uninformed players. Since the number of informed players in this phase is guaranteed to be larger than $n - \sqrt{n}(\ln n)^4$, each uninformed player has probability at least

$$\frac{n-\sqrt{n}(\ln n)^4}{n} \ = \ 1 - \frac{(\ln n)^4}{\sqrt{n}}$$

to receive a rumor due to a pull transmission in each round of this phase. Consequently, we need only a constant number of rounds until all players are informed, w.h.p.

The exponential-growth phase takes $\log_3 n \pm O(\ln\ln n)$ rounds. During this phase the number of transmissions grows exponentially from round to round. Therefore, we send only $O(n)$ messages during this phase. All other phases have length only $O(\ln\ln n)$. Thus, even if we assume $2n$ transmissions in each of these rounds, the total number of transmissions is only $O(n\ln\ln n)$. This completes the proof of Theorem 2.1.

□

3 The median-counter algorithm

The push&pull algorithm relies heavily on a very exact estimation of the expansion of the set of informed players. The algorithm has to be executed for exactly $\log_3 n + \Theta(\ln\ln n)$ rounds. For example, a constant fraction of players remains uninformed if the algorithm terminates after $(1-\epsilon)\log_3 n$ rounds, and the algorithm uses $\Theta(n\ln n)$ transmissions when terminating after $(1+\epsilon)\log_3 n$ rounds, for any constant $\epsilon > 0$. A robust algorithm requires a more flexible, distributed termination mechanism that recognizes when all players are informed. This termination mechanism is described in the following.

Let r denote the rumor being considered. During the course of the algorithm each player v can be in one out of four states A, B, C, or D (with respect to r). State A means the player has not yet received the rumor. In all other states, the player knows the rumor. When a player is in one of the states B or C it pushes and pulls the rumor r along every established connection. In state D the player does not propagate the rumor anymore. Each player in state B holds a counter $\mathtt{ctr}(v,r)$. We say a player v is in state B-m if $\mathtt{ctr}(v,r) = m$. These counters are irrelevant in other states. The transitions between different states are defined as follows.

- State A: The player v does not know r. (For the purpose of analysis, we assume that $\mathtt{ctr}(v,r) = 0$ in this state.) If a player v in state A receives r only from players in state B then it switches to state B-1. If a player in state A receives r from a player in state C then it switches to state C.

- State B-m: The player v knows r and $\mathtt{ctr}(v,r) = m$. (The player injecting the rumor starts in state B-1.)

  Median rule: If during a round a player v in state B-m receives r from more players in state B-m' with $m' \geq m$ than from players in state A and B-m'' with $m'' < m$ then it switches to state B-$(m+1)$, i.e., increases its counter.

  There is one exception to this rule. If $\mathtt{ctr}(v,r)$ is increased to $\mathtt{ctr}_{\max}$ (where $\mathtt{ctr}_{\max} = O(\ln\ln n)$ is a suitable integer) then v switches to state C. Furthermore, if a player in state B receives the rumor from a player in state C then it switches to state C, too.

- State C: Every player stays in this phase for at most $O(\ln\ln n)$ rounds, and then switches to state D, i.e., it terminates the rumor spreading.

Roughly speaking, the counters in state B are used in order to determine the point in time when the algorithm switches from the exponential-growth phase into the quadratic-shrinking phase. A counter value of $\mathtt{ctr}_{\max}$ indicates that $n/\mathrm{polylog}(n)$ players are informed so that it is sufficient to continue the propagation for only $O(\ln\ln n)$ rounds (which is done in state C). In order to make sure that the median-counter algorithm terminates even in the very unlikely event that the counter mechanism fails, we determine that every player stops propagating the rumor after some fixed number of $O(\ln n)$ rounds, regardless of its current state.

We investigate the robustness of the median-counter algorithm against different sources of errors and inaccuracies.

- First, we assume the random connections in each round are established using an arbitrary (possibly non-uniform) probability distribution $\mathcal{D} : V \to [0,1]$.

- Second, we assume that an oblivious adversary can specify up to F node failures occurring during the execution of the algorithm. The adversary specifies a set $\mathcal{F}$ of players (not containing the player starting the rumor) that fail to exchange information in some of the rounds (as specified by the adversary). We assume $|\mathcal{F}| \leq F$ and $n\sum_{v\in\mathcal{F}} \mathcal{D}(v) \leq F$.

Clearly, we cannot hope to inform all players when allowing adversarial node failures. Therefore, we are satisfied if the algorithm informs all but $O(F)$ players. (Alternatively, one may assume stochastic rather than adversarial failures, e.g., each random phone call fails with probability F/n. In this case, staying for $\tau = \Theta(\ln\ln n + \ln_{n/F} F)$ rounds in stage C ensures that all players are informed within $O(\ln n + \tau)$ rounds using $O(\tau n)$ transmissions, w.h.p.)

Theorem 3.1 *Assuming an arbitrary distribution $\mathcal{D}$ and up to F node failures as described above, the median-counter algorithms informs all but $O(F)$ players in $O(\ln n)$ rounds using $O(n \ln\ln n)$ transmissions, w.h.p.*

Proof. First we investigate the errorless case. Let w_i be the probability that a player calls player i, let S_t, s_t, U_t, and u_t be defined as above and let g_t be the weight of all informed players: $g_t := \sum_{i \in S_t} w_i$. Consider the following three phases.

Startup: We want to ensure that at least $s_t \in \Omega(\ln n)$ informed players with weight $g_t \geq \frac{\log n}{n}$ are established.

A straightforward analysis shows that $\Theta(\log\log n)$ rounds of push communication suffice to achieve this, w.h.p.. Then, $\Theta(\log\log n)$ rounds of pull-communication establish the wanted number of informed players w.h.p.

Exponential growth: This phase ends when the weight g_t is greater than $\frac{1}{\log n}$.

In this phase the weight h_t of the set of uninformed players H_t with larger weight than $\frac{1}{s_t}$ is of particular interest:

$$h_t := \sum_{i \in U_t : w_i \geq 1/s_t} w_i \ .$$

Note that $|H_t| \leq s_t$ and that the probability of a member of H_t being called by an informed player in S_t is larger than the constant $1 - 1/e$. Therefore, push operations cause an increase of the weight of informed players by the amount of $(1-\epsilon)(1-1/e)h_t$ for some constant $\epsilon > 0$ w.h.p.

In $U_t \setminus H_t$ the fraction of which get only one call in this round is at least $1/e - \epsilon$ for an arbitrary small constant $\epsilon > 0$ w.h.p. The probability that one of these players gets the rumor pushed from S_t is $\frac{s_t}{n}$. The expected number of informed players in the next round is therefore

$$\begin{aligned} \mathbf{E}[s_{t+1}] &\geq s_t + \frac{s_t}{n}(1/e - \epsilon)(n - s_t - |H_t|) \\ &\geq s_t(1 + (1/e - \epsilon)(1 - \frac{2s_t}{n})) \ . \end{aligned}$$

If $s_t \leq \frac{n}{\log n}$ for $h_t \leq \frac{1}{2}$ this implies $s_{t+1} \geq s_t(1 + \frac{1}{e} - \epsilon')$ and in the other case $g_{t+1} \geq g_t(\frac{3}{2} - \frac{1}{2e} - \epsilon')$ for some arbitrary small $\epsilon' > 0$.

So after some $O(\log n)$ rounds we have either $g_t \geq \frac{n}{\log n}$ or $s_t \geq \frac{2n}{\log n}$. In the second case every player with weight larger than $\frac{c \log^2 n}{n}$ is informed in the next round w.h.p. Furthermore, the expected weight of all informed players is $\mathbf{E}[g_{t+1}] \geq \sum_{i=1}^n w_i^2 s_t$. It turns out that this sum is minimal for the uniform probability distribution. Hence, $\mathbf{E}[g_{t+1}] \geq \frac{s_t}{n}$. Because the weights are upper-bounded we can apply Chernoff bounds and get $g_{t+1} \geq \frac{s_t}{2n} \geq \frac{1}{\log n}$.

For the number of messages note that in all but one round $s_t \leq \frac{2n}{\log n}$. Therefore, the number of messages is bounded by $O(n)$.

Now we discuss how often a counter of a player will be increased during this phase. We consider a player i with weight w_i who is informed during this phase.

1. $w_i \geq \frac{3 \log n}{n}$

   In every round at least $2 \log n$ uninformed players call i, while i receives a call only from at most $\log n$ informed players ($s_t \leq \frac{2n}{\log n}$); i's push call can be neglected. So, this player will communicate with more uninformed than informed players in each round and the median rule prevents an increment of i's counter.

2. $w_i \leq \frac{3 \log n}{n}$

   We allow that during the time interval $\{a, \ldots, b\}$ for which we have $\frac{1}{\log^2 n} \leq w_i s_t \leq c \log n$ the counter of P_i is increased in every round t.

   In every round g_t or s_t (but possibly not both) grows by a factor $\alpha > 1$. Nevertheless they interact pairwise, since the expected number of uninformed nodes informed by a pull is $u_t g_t$. Therefore we have $s_{t+1} \geq (1-\epsilon) u_t g_t \geq n g_t (1 - \epsilon')$ for $\epsilon, \epsilon' > 0$ w.h.p. On the other hand, every informed node pushes in every round such that $g_{t+1} \geq \frac{s_t}{2n \log n}$ w.h.p. So, this time interval is bounded by $O(\log\log n)$.

   At any time step after b the number of uninformed players calling P_i is higher than the number of informed players calling P_i for the same reasons as in 1.

   At every round t before a we concentrate on weights w_i with $w_i \leq \frac{1}{s_t \log^2 n}$. The probability that a player with such a weight is called by an informed player is smaller than $1 - (1 - \frac{1}{s_t \log^2 n})^{s_t} \leq \frac{1}{\log^2 n}$. Let q_i be the number of the players which increase their counter at least i times before round a and let $q_0 = s_t$. In the worst case all players stay in this situation for the whole phase. Only q_i players can cause an increase for a counter larger than i. The probability that such a player calls another is $\frac{q_i}{s_t \log^2 n}$. Therefore, we have $\mathbf{E}[q_{i+1}] \leq \frac{q_i^2}{s_t \log^2 n}$. It follows $\frac{q_{i+1}}{s_t} \leq c \frac{q_i^2}{s_t^2 \log^2 n}$ if $q_i \in \Omega(\log n)$; and if $q_i \leq O(\log n)$, then

$q_{i+c'} = 0$ for some constants c, c' w.h.p. This proves $q_{O(\log\log n)} = 0$. So, there are no players whose counters will be increased more than some $c \log\log n$ time during this phase.

Quadratic-shrinking: This phase ends, when all players have left states A or B.

The probability for each uninformed player to remain uninformed is at most $1 - g_t$, if we consider only pull-communication. Therefore we have $\mathbf{E}[u_{t+1}] \leq u_t(1-g_t)$, which implies

$$u_{t+1} \leq u_t(1-g_t)\left(1 + \frac{\log n}{\sqrt{n}}\right) \quad \text{w.h.p.}$$

The expected weight of the uninformed player of the next round is $\mathbf{E}[1 - g_{t+1}] = (1-g_t)^2$. Note that $\max_{i \in U_t} w_i \leq \frac{c \log^2 n}{n}$. Therefore, applying Chernoff bounds it follows that

$$1 - g_{t+1} \leq (1-g_t)^2\left(1 + \frac{\log^2 n}{\sqrt{n}}\right) \quad \text{w.h.p.}$$

It is clear that after some $O(\log\log n)$ rounds we have $1 - g_{t+1} \leq \frac{2\log^2 n}{\sqrt{n}}$. Then, some constant number of rounds of pull will sufficiently decrease the probability of an uninformed player remaining in state A.

Since in every round each counter may be incremented only once, it suffices to choose $\mathtt{ctr}_{\max} \geq c \log\log n$ for some constant c independent of $\mathcal{D}$.

It remains to show that after some additional $O(\log\log n)$ rounds all counters reach $\mathtt{ctr}_{\max}$. Consider the time point at which all players are informed. Clearly, all counters are at least 1. Then, in every step i each counter is at least $i+1$. Therefore the distributional algorithm ends after $O(\log\log n)$ rounds.

Since every player produces only one random call in each round the overall number of messages in this phase is bounded by $O(n \log\log n)$.

Now we focus on the case of $F \leq \frac{1}{4}n$ node failures with weight F/n. We assume that if a node failure occurs on v that v terminates, i.e. switches to state D without learning the rumor. The analysis of the startup and exponential phases can be easily adapted to this case, since the growth of informed nodes proceeds more slowly but still exponentially. We now investigate the situation in the double exponential shrinkage phase.

Let $\mathcal{F}$ be the set of nodes which may be disconnected in some rounds. Then S_t and U_t are defined as the set of informed and uninformed nodes, excluding the nodes in $\mathcal{F}$; u_t, s_t, and g_t are defined as before. The probability that a node remains uninformed is at most $1-g_t$ per round. Therefore we can conclude that w.h.p. $u_{t+1} \leq (1-g_t)u_t$. Similarly to the error-free case, we can conclude that $1 - g_{t+1} \leq \frac{F}{n} + (1 + \frac{\log^2 n}{n})(1-g_t)^2$ w.h.p. This recursion converges in $O(\log\log n)$ rounds to $1 - g_{t'} \in O(\frac{F}{n})$. This implies a maximum number of $O(F)$ uninformed nodes within the next round.

The main problem for the error case is to verify that the number of messages does not exceed $O(n \log\log n)$. We prove this by showing that at least $O(n/\log n)$ players have reached state C or D, by the time the first error-free players reach state D. The remaining error-free players can only cause $O(\log n)$ messages each, where F faulty players do not add further messages. We start our analysis at the moment when only $F' \in O(F)$ nodes with weight F'/n remained uninformed. Let us assume that all informed players are in the state B-1.

Let $Z_{t,m}$ be the set and $y_{t,m}$ the weight of error-free nodes in round t with $\mathtt{ctr}(v) = m$. The probability that a node in $Z_{t,m}$ is increased is at least $\sum_{i=m}^{\mathtt{ctr}_{\max}} y_{t,i}$. We want to prove that in the triangular section where $t \leq km$ for some constant k, $y_{t,m}$ decreases exponentially in t. For the analysis we allow that some of the counters may be decreased. The aim of this modification is that the series $y_{t,1}, \ldots, y_{t,m_t}$ is exponentially increasing, the series $y_{m_t,t}, y_{m_t+1,t}, \ldots$ is exponentially decreasing, and the weight $y_{t,m_t+1} \geq \frac{1}{2}$ contains the rest of the weight. More formally, $\forall i \leq m_t \;:\; y_{t,i} \leq \alpha y_{t,i+1}$ and $y_{t,m_t+1} = 1 - F'/n - \sum_{i=0}^{m_t} y_{t,i}$ for some $\alpha > 1$.

By decreasing some of the counters it can be ensured that in the next round we have $\forall i \leq m_t \;:\; y_{t,i} \leq \alpha\, y_{t,i+1}$ and $y_{t+1,i} \leq \frac{1+\alpha}{2}\, y_{t,i}$. This follows by the fact that $\sum_{i=j}^{m_t} y_{t,j} \geq \frac{1}{2}$ and by reducing the number of players increasing their counter to a fraction of $\frac{1}{2}$ each. After some constant number of rounds c we have $y_{t+c,m_t+1} \geq \alpha y_{t+c,m_t}$. Then, we increase $m_{t+c} := m_t + 1$ and get the claimed triangular section.

Therefore, after some $O(\log\log n)$ rounds only a fraction of $O(n/\log n)$ players has a smaller counter than $c \log\log n$. □

4 Lower bound for address-oblivious algorithms

Our first lower bound shows that the two presented push&pull algorithms achieve the best possible results for the class of address-oblivious algorithms. Clearly, any algorithm requires $\Omega(\ln n)$ rounds in order to inform all players. In addition, we show that any address-oblivious algorithm requires $\Omega(n \ln\ln n)$ transmissions, regardless of the number of rounds. We assume the random phone call model using the uniform distribution.

Theorem 4.1 *Any address-oblivious algorithm guaranteeing that all but a fraction f of the players receive the ru-*

mor with constant probability needs to perform $\Omega(n \ln\ln \frac{1}{f})$ *transmissions in expectation.*

Proof. Let us fix an address-oblivious algorithm $\mathcal{A}$. Depending on the execution of $\mathcal{A}$, we will partition the rounds into contiguous phases such that the number of transmissions during the first i phases is at least $(i-1)n/4 = \Omega(in)$. (The actual length of the phases depends possibly on the outcome of random experiments influencing the execution of $\mathcal{A}$. Thus, the length of the phases might give some evidence about the outcome of some random experiments. The following statement, however, holds regardless of this evidence.) Let U_i denote the number of uninformed players at the end of phase i, and define $u(i) = n\exp(-2^i + \frac{3}{2})$. We will show by induction that $U_i \geq u(i)$, w.h.p. Consequently, $\mathcal{A}$ needs $\Omega(\ln\ln \frac{1}{f})$ phases and, hence, $\Omega(n \ln\ln f)$ transmissions in order to inform all but a fraction f of the players. Clearly this yields the Theorem.

Phases are defined as follows. Phase 1 starts with the round in which the rumor is generated. If phase i ends in round t then phase $i+1$ starts in round $t+1$. Thus, it remains to describe when a phase ends. We distinguish sparse and dense phases. A *sparse phase* contains at most $n/2$ transmissions. The length of these phases is maximized, that is, a sparse phase ends in round t if adding round $t+1$ to the phase would result in more than $n/2$ transmissions. A *dense phase* consists of only one round containing more than $n/2$ transmissions. Observe that the number of transmissions during the phases 0 to i is at least $(i-1)n/4$ because any pair of consecutive phases contains at least $n/2$ transmissions by construction.

Now assume by induction that the number of uninformed players at the beginning of phase i is at least $u(i-1)$. We have to show that the number of uninformed players at the end of phase i is at least $u(i)$, w.h.p.

For $1 \leq k \leq u(i-1)$, let x_k denote a 0-1 random variable indicating whether the kth of those players that are uninformed at the beginning of round i receives a message containing the rumor during the round. We claim

$$\mathbf{Pr}[x_k = 0] \geq \frac{u(i-1)}{\mathrm{e}n} \ .$$

The arguments leading to this inequality are different for sparse and dense rounds.

- Suppose phase i is sparse. Then $\mathcal{A}$ sends at most $\frac{n}{2}$ messages during this phase. Each of these messages is initiated without knowing the receiver because decisions are placed in an address-oblivious fashion. As connections are chosen uniformly at random, the probability that any particular message reaches player k is $\frac{1}{n}$. Consequently, $\mathbf{Pr}[x_k = 1] \leq \frac{n}{2} \cdot \frac{1}{n} \leq \frac{1}{2}$ so that $\mathbf{Pr}[x_k = 0] \geq \frac{1}{2} \geq \frac{u(i-1)}{\mathrm{e}n}$.
- Now suppose phase i is dense. Then the phase consists of only one round. In this case, the probability p_1 that player k does not call an informed player is at least $\frac{u(i-1)}{n}$. Furthermore, the probability p_2 that player k is not called by any other player is at least $\frac{1}{\mathrm{e}}$. As these two probabilities are independent, $\mathbf{Pr}[x_k = 0] = p_1 p_2 \geq \frac{u(i-1)}{\mathrm{e}n}$.

Since $U(i) = \sum_{k=1}^{u(i-1)} (1 - x_k)$, we obtain

$$\begin{aligned}
\mathbf{E}[U(i)] &= \sum_{k=1}^{u(i-1)} \mathbf{Pr}[x_k = 0] \\
&\geq \frac{u(i-1)^2}{\mathrm{e}n} = \frac{(n\exp(-2^{i-1} + \frac{3}{2}))^2}{\mathrm{e}n} \\
&= n\exp(-2^i + 2) = \sqrt{\mathrm{e}}u(i) \ .
\end{aligned}$$

In particular, $u(i) \leq (1 - \frac{1}{3})\mathbf{E}[U(i)]$. Observe that the random variables x_k are slightly dependent since the random interconnections used for transmissions in phase i form partial permutations on the caller sites. This dependence, however, is negative so that we can apply a Chernoff bound [3]. Assuming $u(i) \geq (\ln n)^2$, we obtain

$$\begin{aligned}
\mathbf{Pr}[U_i < u(i)] &\leq \mathbf{Pr}\left[U_i < (1 - \tfrac{1}{3})\mathbf{E}[U(i)]\right] \\
&\leq \exp(-\tfrac{1}{18}\mathbf{E}[U(i)]) \\
&\leq \exp(-\tfrac{1}{12}u(i)) = O(n^{-\alpha}) \ ,
\end{aligned}$$

for any positive constant α. This completes the proof of Theorem 4.1. ☐

5 Lower bound for general algorithms

The above lower bound for address-oblivious algorithms does not hold for those rumor spreading algorithms that can base their decisions on the addresses of communication partners. In the introduction, we give an example showing how all players can be informed using only $O(n)$ transmissions. This unrealistic algorithm, however, requires $\Theta(n \ln n)$ rounds. Now we investigate whether there is an algorithm that is both time-optimal (i.e., using only $O(\log n)$ rounds) and communication-optimal (i.e., using only $O(n)$ transmissions) The following lower bound answers this question negatively. Again, we assume the random phone call model using the uniform distribution.

Theorem 5.1 *Any distributed rumor spreading algorithm guaranteeing that all but a fraction* $o(1)$ *of the players receive the rumor within* $O(\ln n)$ *rounds with constant probability needs to perform* $\omega(n)$ *transmissions in expectation.*

Proof. The difficulty in analyzing arbitrary distributed rumor spreading algorithms is that the distribution of the rumor can be a highly dependent process although the underlying random calling mechanism is generated by n independent experiments in each round. For example, if player

1 is the only player with an odd address sending the rumor to players with even addresses then the success of the algorithm is highly dependent on the event that player 1 receives the rumor. This small example (not even involving any additional communication) shows that the analysis needs more than simply applying martingales or Chernoff bounds.

Our basic trick in the following analysis is that we choose a random sample of the players that can be guaranteed to act independently. This independence, however, can be guaranteed only for about $\frac{1}{8}\log n$ rounds. Of course, this number of rounds is not enough to inform all players about a rumor initiated by a single player. Therefore, let us assume for the time being that the rumor is spread already to at least half of the players and we consider the next $T = \lfloor \frac{1}{8}\log n \rfloor$ rounds.

Consider an arbitrary rumor spreading algorithm $\mathcal{A}$. Let $U_V \le n/2$ denote the number of initially uninformed players. (In order to be able to extend our result to more than T rounds later, we assume that the initially uninformed players are known by all players in the system, e.g., assume that $\{1, \ldots, U_V\}$ is the set of initially uninformed players.) Let X_V denote a random variable describing the number of messages sent during the T rounds. Furthermore, let U'_V denote a random variable describing the number of uninformed players after round T. (These random variables are with respect to the random phone calls and any kind of other random decisions made by $\mathcal{A}$.)

Let S denote a set of $m = \lfloor n^{1/8} \rfloor$ players chosen randomly from V. The set S will be our random sample. Let U_S denote the random variable describing the number of initially uninformed players in S (with respect to the random choice of S.) Let X_S denote a random variable describing the number of messages received by the players in S, and let U'_S denote the random variable describing the number of uninformed players in the set S after the last round. (These random variables are with respect to random decisions of $\mathcal{A}$ and the random choice of S.)

Recall that the communication graph G_t in round t is obtained by a distributed random process, i.e., each player v chooses a player u from V at random and v calls u. This random process generates a probability distribution $\mathcal{D}$ on the set $\mathcal{G}$ of possible communication graphs. Repeating this random process for T rounds extends the probability distribution $\mathcal{D}$ to the sample space $\mathcal{G}^T$.

In many parts of the following analysis, we will assume a slightly different probability distribution $\mathcal{D}'$ on $\mathcal{G}$ that is easier to handle than $\mathcal{D}$. Instead of letting each player call a random other player, we assume that the connections are established as follows. In each round t,

- we choose uniformly at random a collection of m disjoint subsets $B_t(v)$ ($v \in S$), each containing m players from $V \setminus S$; (once these sets are chosen, the players in S can act fully independently)
- each player $v \in S$, chooses at random an integer $\delta(v) \ge 0$ with $\mathbf{Pr}[\delta(v) = i] = \dfrac{1}{\mathrm{e}i!}$; (in the very unlikely case that $\delta(v) \ge m$, set $\delta(v) = m - 1$)
- each player $v \in S$, chooses independently and uniformly at random a set of $\delta(v) + 1$ different players $u_0(v), \ldots, u_{\delta(v)}(v)$ from $B_t(v)$.

We determine that every player $v \in S$ calls player $u_0(v)$, and the players $u_1(v), \ldots, u_{\delta(v)}(v)$ call v. Every player for which we have not yet specified whom to call simply chooses a communication partner from $V \setminus S$ independently and uniformly at random. Clearly, $\mathcal{D}$ and $\mathcal{D}'$ are different distributions. The following lemma, however, shows that these distributions are closely related.

Lemma 5.2 *The total variation distance between $\mathcal{D}$ and $\mathcal{D}'$ on $\mathcal{G}^T$ is $O(n^{-1/4})$.*

Based on this bound, we are able to give the following lemma comparing the behavior of the complete system $V|\mathcal{D}$ with that of the small system $S|\mathcal{D}'$.

Lemma 5.3 *For $\beta \ge 0$, $u \ge n^{-1/16}$, $0 \le p \le 1$,*

a) $\mathbf{E}[X_V|\mathcal{D}] \le \beta n \Rightarrow$
$\mathbf{Pr}\left[X_S > \frac{\beta m}{p}|\mathcal{D}'\right] \le p + O(n^{-1/4})$,

b) $U_V \ge un \Rightarrow$
$\mathbf{Pr}\left[U_S < \frac{um}{2}\right] = O(n^{-1/4})$, *and*

c) $\mathbf{Pr}[U'_S \ge um|\mathcal{D}'] < p \Rightarrow$
$\mathbf{Pr}\left[U'_V < \frac{un}{2}|\mathcal{D}\right] \le p + O(n^{-1/4})$.

Informally, this lemma states that it is sufficient to analyze $S|\mathcal{D}'$ in order to estimate $V|\mathcal{D}$. In fact, restricting to the smaller and simpler system $S|\mathcal{D}'$ will enable us to deal with the complex dependencies in the original system $V|\mathcal{D}$. The following lemma summarizes our analysis for $S|\mathcal{D}'$.

Lemma 5.4 *Let c denote a suitable constant. Suppose $U_S \ge m/\alpha$ and $X_S \le \beta m$ with $\alpha \ge 4$ and $\beta \ge 1$. Then*

$$U'_S \ge m\alpha^{-\exp(c\alpha\beta)} \quad ,$$

with probability $1 - O(n^{-1/4})$, provided that α, β, c are not too large so that $\alpha^{-\exp(c\alpha\beta)} \ge n^{-1/16}$.

(Due to space limitations, we omit the proof of all three lemmas. Lemma 5.2 and 5.3 can be shown using standard methods from probability theory like, e.g., Chernoff bounds and the Markov inequality. The proof of Lemma 5.4 is more interesting. We transform the probabilistic system $S|\mathcal{D}'$ into a deterministic token game, which then can be analyzed combinatorially.)

Combining Lemma 5.3 and 5.4, we obtain the following result for $V|\mathcal{D}$. Suppose $U_V \leq n/\alpha$ and $\mathbf{E}[X_V \leq \beta n]$ with $2 \leq \alpha \leq n^{1/16}$ and $\beta \geq 0$. Applying Lemma 5.3 a) and b) yields

$$X_S \leq \kappa\beta m \quad \text{and} \quad U_S \geq \frac{m}{2\alpha} \ ,$$

with probability at least $1 - \frac{1}{\kappa} - O(n^{-1/4})$, for any $\kappa > 0$. Now applying Lemma 5.4 yields

$$U'_S \quad \geq \quad m\alpha^{-\exp(c\alpha(\kappa\beta+1))} \ ,$$

with probability $1 - \frac{1}{\kappa} - O(n^{-1/4})$. Finally, we can conclude from Lemma 5.3 c) that

$$U'_V \quad \geq \quad \frac{n}{2}\alpha^{-\exp(c\alpha(\kappa\beta+1))} \ , \qquad (1)$$

with probability $1 - \frac{1}{\kappa} - O(n^{-1/4})$. Assuming $n \gg \kappa$, this probability is lower-bounded by $1 - \frac{2}{\kappa}$.

For the time being, let us assume α and β are constants. Then equation 1 can be interpreted as follows. Starting with n/α uninformed players (possibly known by all players), performing $X_v \leq \beta n$ transmissions in $\lceil \frac{1}{8} \log\log n \rceil$ rounds reduces the number of uninformed players only by some constant factor, with probability at least $1 - \frac{2}{\kappa}$. Now let us consider the execution of c phases of length at most $\lceil \frac{1}{8} \log\log n \rceil$ each, for any constant $c \geq 1$. Suppose we spend at most βn transmissions in each of these phases. Then the number of uninformed players after all c phases is $\Theta(n)$, with probability $1 - \frac{2c}{\kappa}$. Let us set $\kappa \geq 2c/\epsilon$, for any constant $\epsilon > 0$. Then spending $\Theta(n)$ transmissions in $O(\ln n)$ rounds leaves $\Theta(n)$ uninformed players, with probability $1 - \epsilon$. (A rigorous analysis based on inequality 1 shows that informing all but a fraction f of the players with constant probability requires $\mathbf{E}[X_V] = \Omega(\ln^{[2k]} \frac{1}{f})$, where $\ln^{[x]}$ denotes the natural logarithm iterated for x times.) Hence, Theorem 5.1 is shown. $\square$

References

[1] D. Agrawal, A. El. Abbadi, R. C. Steinke. Epidemic Algorithms in Replicated Databases. In *Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 161-172, 1997.

[2] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic Algorithms for Replicated Database Maintenance. In *Proceedings of the 6th ACM Symposium on Principles of Distributed Computing*, pages 1-12, 1987.

[3] D. Dubhashi, D. Ranjan. Balls and Bins: A Study in Negative Dependence. In *Random Structures & Algorithms*, 13(2):99-124, 1998.

[4] L. Gasieniec, A. Pelc, Adaptive broadcasting with faulty nodes, In *Parallel Comput. 22 (1996) 903-912., 1996.*

[5] S. Hedetniemi, S. Hedetniemi, and A. Liestman. A Survey of Gossiping and Broadcasting in Communication Networks. Networks 18, 1988, 319-349.

[6] J. Hromkovic, R. Klasing, B. Monien, and R. Peine. Dissemination of Information in Interconnection Networks (Broadcasting & Gossiping). In *Combinatorial Network Theory*, pp. 125–212, D.-Z. Du and D.F. Hsu (Eds.), Kluwer Academic Publishers, Netherlands, 1996.

[7] R. Golding, D. Long. Accessing Replicated Data in a Large-Scale Distributed System. UCSC-CRL-91-01, Santa Cruz CA, 1991.

[8] R. Guy, G. Popek, T. Page, Jr. Consistency Algorithms for Optimistic Replication. In *Proceedings of the First International Conference on Network Protocols*, IEEE, 1993.

[9] R. Ladin, B. Liskov, L. Shrira, S. Ghemawat. Providing high availability using lazy replication. In *ACM Transaction on Computer Systems*, 10(4):360, 1992.

[10] T. Leighton, B. Maggs, R. Sitamaran. On the fault tolerance of some popular bounded-degree networks. In *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*, pages 542-552, 1992.

[11] D. Malkhi, Y. Mansour, M. K. Reiter. On Diffusing Updates in a Byzantine Environment. In *Proceedings of the 18th IEEE Symposium on Reliable Distributed Systems*, pages 134-143, 1999.

[12] M. Rabinovich, N. Gehani, A. Kononov. Scalable update propagation in epidemic replicated databases. AT&T Bell Labs Technical Memorandum 112580951213-11TM, 1995.