

Rumor Spreading with Bounded In-Degree*

Sebastian Daum, Fabian Kuhn, and Yannic Maus

Department of Computer Science, University of Freiburg
79110 Freiburg, Germany

daum.sebastian@gmail.com, {kuhn, yannic.maus}@cs.uni-freiburg.de

Abstract

We consider a variant of the well-studied gossip-based model of communication for disseminating information in a network, usually represented by a graph. Classically, in each time unit, every node u is allowed to contact a single random neighbor v . If u knows the data (rumor) to be disseminated, node v learns it (known as *push*) and if node v knows the rumor, u learns it (known as *pull*). While in the classic gossip model, each node is only allowed to contact a single neighbor in each time unit, each node can possibly be contacted by many neighboring nodes. If, for example, several nodes pull from the same common neighbor v , v manages to inform all these nodes in a single time unit.

In the present paper, we consider a restricted model where at each node only one incoming request can be served in one time unit. As long as only a single piece of information needs to be disseminated, this does not make a difference for push requests. It however has a significant effect on pull requests. If several nodes try to pull the information from the same common neighbor, only one of the requests can be served. In the paper, we therefore concentrate on this weaker pull version, which we call *restricted pull*.

We distinguish two versions of the restricted pull protocol depending on whether the request to be served among a set of pull requests at a given node is chosen adversarially or uniformly at random. As a first result, we prove an exponential separation between the two variants. We show that there are instances where if an adversary picks the request to be served, the restricted pull protocol requires a polynomial number of rounds whereas if the winning request is chosen uniformly at random, the restricted pull protocol only requires a polylogarithmic number of rounds to inform the whole network. Further, as the main technical contribution, we show that if the request to be served is chosen randomly, the slowdown of using restricted pull versus using the classic pull protocol can w.h.p. be upper bounded by $O(\Delta/\delta \cdot \log n)$, where Δ and δ are the largest and smallest degree of the network.

1998 ACM Subject Classification G.2.2 Graph Theory

Keywords and phrases Rumor Spreading, Push-Pull, Gossiping, Coupling

Digital Object Identifier [10.4230/LIPIcs.xxx.yyy.p](https://doi.org/10.4230/LIPIcs.xxx.yyy.p)

1 Introduction

Gossip-based communication models have received a lot of attention as a simple, fault-tolerant, and in particular also scalable way to communicate and disseminate information in large networks. The classic application of gossip-based network protocols is the spreading of information in the network, specifically the problem of broadcasting a single piece of information to all nodes of a network, in this context also often known as *rumor spreading*,

* Research supported by ERC Grant No. 336495 (ACDC)



© S. Daum, F. Kuhn and Y. Maus;
licensed under Creative Commons License CC-BY

23rd International Colloquium on Structural Information and Communication Complexity.



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

e.g., [6, 8, 10, 5, 12, 16]. On top of this, gossip-based protocols have for example also been proposed for applications such as maintaining consistency in a distributed database [6], for data aggregation problems [17, 18, 4], or even to run arbitrary distributed computations [2].

The best studied gossip strategy is the *random phone call model*, which was first considered in [10]. We are given a network graph $G = (V, E)$ where initially a source node $s \in V$ knows some piece of information (*rumor*) and the objective is to disseminate the rumor to all nodes of G . Typically, time is divided into synchronized rounds, where in each round, every node can contact a random neighbor and if u contacts v , an interaction between u and v is initiated for the current round. For spreading a rumor, two basic modes of operation are distinguished. Nodes that already know the rumor can PUSH the information to the randomly chosen neighbor [10] or nodes that do not yet know the rumor can PULL the information from the randomly chosen neighbor [6]. In much of the classic work, the network G is assumed to be a complete graph. In that case, it is not hard to see that PUSH and PULL both succeed in $O(\log n)$ rounds and that the total number of interactions of each node can also be bounded by $O(\log n)$. In [16], it is shown that when combining PUSH and PULL (in the following referred to as PUSH-PULL), the average number of interactions per node is only $\Theta(\log \log n)$.

Mostly in recent years, PUSH, PULL, and PUSH-PULL have also been studied for more general network topologies, e.g., [3, 8, 9, 5, 12, 11, 13], with [5, 12] and [11, 13] studying the time complexity as a function of the graph's conductance and vertex expansion, respectively. E.g., in [12], it is shown that with high probability (w.h.p.), the running time of PUSH-PULL can be upper bounded by $O((\log n)/\phi(G))$, where n is the number of nodes and $\phi(G)$ is the conductance of the network graph G .

While in gossip protocols, each node can initiate at most one interaction with some neighbor, even if each node contacts a uniformly random neighbor, the number of interactions a node needs to participate in each round can be quite large. In complete graphs and more generally in regular graphs, the total number of interactions per node and round can easily be upper bounded by $O(\log n)$. However in general topologies a single node might be contacted by up to $\Theta(n)$ neighboring nodes. As an extreme case, consider a star network where a single center node is connected to $n - 1$ leaf nodes. Even if the rumor initially starts at a leaf node, PUSH-PULL manages to disseminate the rumor to all nodes in only 2 rounds. Clearly, in these 2 rounds, the center node has to interact with all $n - 1$ leaf nodes. In fact, all recent papers which study the time complexity of the random PUSH-PULL protocol critically rely on the fact that a node can be contacted by many nodes in a single round, e.g., [12]. In some cases, this behavior might limit the implementability and thus the applicability of the proven results for this gossip protocol. In order to obtain scalable systems, ideally, we would like to not only limit the number of interactions each node initiates, but also the number of interactions each node participates in.

In the present paper, we therefore study a weaker variant of the described random gossip algorithms. In each round, every node can still initiate a connection to one uniformly random neighbor. However, if a single node receives several connection requests, only one of these connections is actually established. When disseminating a rumor by using the PUSH protocol, this restriction does not limit the progress of the algorithm. In a given round, a node v learns the rumor if and only if at least one PUSH request arrives at v . However, when using the PULL protocol, the restriction can have a drastic effect. If a node v receives several PULL requests from several nodes that still need to learn the rumor, only one of these nodes can actually learn the rumor in the current round. In our paper, we therefore concentrate on the PULL protocol and we define RPULL (*restricted PULL*) as the described weak

variant of the PULL algorithm: In each RPULL round, every node that still needs to learn the rumor contacts a random neighbor. At every node that knows the rumor, one of the incoming requests (if there are any) is selected and the rumor is sent to the corresponding neighbor. By PUSH-RPULL we denote the combination of RPULL with a simultaneous execution of the classic PUSH protocol.

Contributions

We first consider two versions of the RPULL protocol which differ in the way how one of the incoming requests is selected. Assume that in a given round some informed node v receives RPULL requests from a set of neighbors R_v . In the *adversarial* RPULL protocol, an (adaptive) adversary picks some node $u \in R_v$ which will then learn the rumor. In the *random* RPULL protocol, we assume that a uniformly random node $u \in R_v$ learns the rumor (chosen independently for different nodes and rounds). While the choice of which neighbor a node (actively) contacts with a request is under the control of the protocol, it is not necessarily clear how one of the incoming requests in R_v is chosen. If the node can only answer one request per time unit and the requests do not arrive at exactly the same time, the first request might be served and all others dropped. Or even if requests arrive at the same time, it might be the underlying network infrastructure or operating system which picks one request and drops the others. If it is reasonable to assume that the incoming requests are served probabilistically and independently, we believe that random RPULL provides a good model. Otherwise, the adversarial assumption allows to study the worst-case behavior.

As a first result, we prove that the running times of the two RPULL variants are essentially the same on trees. Secondly, we show that there are instances for which there is an exponential gap between the running times of the two RPULL variants. We give an instance where for every source node the random RPULL protocol informs all nodes of the network in polylogarithmic time, w.h.p., whereas, for every source, the adversarial RPULL algorithm requires time $\Omega(\sqrt{n})$ to even succeed with a constant probability.

In the second part of the paper, we have a closer look at the performance of the random RPULL protocol. Consider a graph G and let δ and Δ denote the smallest and largest degree of G . In each round, in expectation, each informed node receives at most Δ/δ requests. Hence, if an uninformed node u sends an RPULL request to an informed node, u should receive the rumor with probability at least $\Omega(\delta/\Delta)$. Consequently, intuitively, the slowdown of using random RPULL instead of the usual PULL protocol should not be more than $\tilde{O}(\Delta/\delta)$.¹ We prove that this intuition is correct. For every given instance, we show that if the PULL algorithm informs all nodes in \mathcal{T} rounds with probability p , for the same instance, the random RPULL algorithm manages to reach all nodes in time $O(\mathcal{T} \cdot \frac{\Delta}{\delta} \cdot \log n)$ with probability $(1 - o(1))p$.² While the statement might seem very intuitive, its formal proof turns out quite involved. Formally, we prove a stronger statement and show that a single round of the PULL protocol is w.h.p. *stochastically dominated* by $O(\frac{\Delta}{\delta} \cdot \log n)$ rounds of random RPULL in the following sense. We give a coupling between the random processes defined by PULL and random RPULL such that for every start configuration, w.h.p., the set of nodes informed after $O(\frac{\Delta}{\delta} \cdot \log n)$ rounds of random RPULL is a superset of the set of nodes informed in a single PULL round. The same holds for simulating one round of PUSH-PULL with PUSH-RPULL. A similar coupling between rumor spreading

¹ Here \tilde{O} hides $\log(n)$ factors.

² Actually, $\frac{\Delta}{\delta}$ can be replaced by $\max_{\{u,v\} \in E} d(u)/d(v)$ in all parts of the paper, where $d(u)$ and $d(v)$ denote the degrees of the nodes.

algorithms has been done in [1] where the authors couple $\log(n)$ rounds of asynchronous with one round of synchronous PUSH-PULL. A coupling between PULL and RPULL in the classic sense, i.e., a coupling which does relinquish the w.h.p. term does not exist. We also show that for such a round-by-round analysis, our bound is tight. That is, there are configurations where $\Omega(\frac{\Delta}{\delta} \log n)$ random RPULL rounds are needed to dominate a single PULL round with high probability.

Notation and Preliminaries

Let $G = (V, E)$ be the n -node network graph. For a node $u \in V$, we use $N(u)$ to denote the set of neighbors of u and $d(u) := |N(u)|$ to denote its degree. Given a set of nodes $S \subseteq V$, we define $N_S(u) := N(u) \cap S$ to be the set of u 's neighbors in S and $d_S(u) := |N_S(u)|$ for the number of neighbors of u in S . The smallest and largest degrees of G are denoted by δ and Δ , respectively. For a set $V' \subseteq V$ we denote with $G[V']$ the graph induced by V' . To indicate a disjoint union of two sets, i.e., $A \cup B$ with $A \cap B = \emptyset$, we write $A \cup B$. For a set of natural numbers $\{1, \dots, k\}$ we only write $[k]$.

When analyzing the progress of an algorithm ALG, the set S_t^{ALG} denotes the set of informed nodes after t rounds and U_t^{ALG} the set of uninformed nodes. When the algorithm is clear from the context we simply write S_t and U_t , furthermore we denote $S = S_0$ and $U = U_0$ for the initial configuration.

2 Separation of Adversarial and Random RPULL

We want to show that the adversarial RPULL can be exponentially slower than the randomized RPULL on general graphs. To show this, we first establish results on the run time of both algorithms on trees. These results might also be of independent interest.

In a tree network let $p_{v,u} = (v = v_0, v_1, v_2, \dots, v_q = u)$ denote the unique path from v to u , though we use that notation also for the set of nodes on that path, i.e., $p_{v,u} = \{v, v_1, v_2, \dots, v_{q-1}, u\}$. For a path p , let $M_p := \sum_{w \in p} d(w)$ be the sum of all degrees on the path.

The next lemma shows that on a tree any form of RPULL is asymptotically as fast as PULL plus an additive term in the order of the degree of the node that initially has the rumor.

► **Lemma 1.** *Let G be a tree network with $S_0 = \{r\}$ and let u be a node in U_0 . Furthermore, let τ be the first round in which $u \in S_\tau$ holds, i.e., the number of rounds until u gets informed.*

- (1) $\mathbf{E}[\tau] = \Theta(M_{p_{r,u}} - d(r))$ for PULL,
- (2) $\mathbf{E}[\tau] = \Omega(M_{p_{r,u}} - d(r))$ for every type of RPULL,
- (3) $\mathbf{E}[\tau] = O(M_{p_{r,u}})$ for adversarial RPULL.

Proof of Lemma 1. (1) We root the tree at the only informed node, r . Note that nodes are not aware of their own parent/child relationships. Consider some time t at which a node r' on the path $p_{r,u}$ is in $S_t \setminus S_{t-1}$, i.e., it just got informed. Thus its child $u' \in p_{r,u}$ on the path is not yet informed, i.e., $u' \in U_t$. In any round $t' \geq t$, in which u' is not informed yet, it requests its parent with probability $1/d(u')$. Thus each uninformed node $u' \in p_{r,u} \setminus \{r\}$ on the path needs $\Theta(d(u'))$ rounds in expectation before it can get informed. Linearity of expectation proves the claim for PULL.

(2) follows from the fact that RPULL is at most as fast as PULL.

(3) For adversarial RPULL divide all rounds $t' \geq t$ in which u' is not yet informed into two types: First rounds in which at least one *sibling* of u' , i.e., the nodes in $N(r') \setminus \{u'\}$,

requests from r' and secondly rounds in which no sibling of u' requests from r' . The first type of rounds is upper bounded by $d(r')$ because every neighbor of r' stops requesting after receiving the rumor. In expectation u' gets the rumor after $d(u')$ rounds of type two; thus in expectation u' is informed within $O(d(r') + d(u'))$ rounds. Applying this recursively to all uninformed nodes on the path $p_{r,u}$, we get the claimed result via linearity of expectation. ◀

► **Lemma 2.** Let G be a tree network with $S_0 = \{r\}$. Then in both random and adversarial RPULL it takes $O(\max_{\text{path } p} M_p + \Delta \log n)$ rounds to fully inform all nodes in V , w.h.p..

Proof Sketch. The full proof appears in Appendix C. Let u be an uninformed node with informed parent v . In each round v either informs one child other than u or u has a probability of $1/d(u)$ to get the rumor. In expectation for each predecessor-successor pair (v, u) on a path p the time for u to get informed is at most $d(v) + d(u)$. A Chernoff bound for geometric random variables (Appendix C, Lemma 20) provides the time bound of $O(M_p + \Delta \log n)$ for that path and a union bound over all paths concludes the proof. ◀

Lemma 2 shows that random RPULL and adversarial RPULL are essentially the same on trees. This does not hold for general graphs.

► **Lemma 3.** There is a graph $G = (V, E)$ of size $\Theta(n)$ with node $r_\alpha \in V$, $d(r_\alpha) \leq 3$, s.t.:

- For $S_0 = \{r_\alpha\}$, w.c.p., the run-time of adversarial RPULL is in $\Omega(\sqrt{n})$.
- For any non-empty $S_0 \subset V$, w.h.p., the run-time of randomized RPULL is in $O(\log^2 n)$.

Proof Sketch. The full proof appears in Appendix C and a figure for visualization appears in Appendix A. We first introduce a special graph type that we call a k -leaf-connected tree (k -LCT). In simple words, a k -LCT is a binary tree with k leaves, but with its k leaves being fully interconnected, i.e., forming a clique. Propagation of the rumor from one node to all nodes of a k -LCT, happens in $O(\log k)$. This also holds true if we embed this k -LCT into a larger graph, as long as the degrees do not grow much in this manner. As long as the k -LCT is uninformed most requests of the leaves will target other leaf nodes, hence it is unlikely for a k -LCT to acquire the rumor through its leaves.

We construct G as follows: We let D_α and D_ζ be two n -LCTs, and we have m l -LCTs that we denote with D_1, D_2, \dots, D_m , where $l := \sqrt{n}$ and $m := c\sqrt{n}$ for some $c > 1$. Their corresponding roots and leaf sets are denoted as $r_\alpha, r_\zeta, r_1, r_2, \dots, r_m$ and $L_\alpha, L_\zeta, L_1, L_2, \dots, L_m$ respectively, and with $l_{X,1}, l_{X,2}, \dots$ we enumerate the leaves of leaf set L_X . Let $C_\alpha = \{c_1, \dots, c_m\}$ be an arbitrary m -sized subset of $D_\alpha \setminus L_\alpha$ and add the following edges:

- Between r and D_ζ : We add one edge from r to r_ζ .
- Between r and D_α : For each $j \in [m \log n]$ we add an edge from r to $l_{\alpha,j}$.
- Between r and D_1, \dots, D_m : For each $i \in [m]$, $j \in [\log n]$ we add an edge from r to $l_{i,j}$.
- Between D_1, \dots, D_m and D_ζ : For each $i \in [m]$ we add an edge from r_i to $l_{\zeta,i}$.
- Between D_1, \dots, D_m and C_α : For each $i \in [m]$ we add an edge from $l_{i,l}$ to c_i .

The idea of the proof is the following: The graph is built in a way that information propagation from D_ζ to the rest of the graph is quick, but not the other way round. In the random RPULL model, wherever the rumor starts, it reaches r quickly and from there r_ζ manages to pull in polylogarithmic time. Then the rumor quickly propagates through D_ζ , and from L_ζ to all LCTs D_1, \dots, D_m and afterwards to D_α .

In the adversarial RPULL model, as long as the rumor starts outside D_ζ , the rumor can quickly spread to r , D_α and a few of the D_i s (i.e., the LCTs D_1, \dots, D_m). But then we let the adversary always prioritize a request at node r from a node in one of the D_i s

over a request from r_ζ to prevent that r_ζ gets the rumor. We can show that the number of informed D_i s grows slowly and hence such requests exist w.h.p. as long as no node in D_ζ is informed. Also, with only few D_i s informed, due to their high degrees, leaf nodes in L_ζ are unlikely to request from a D_i containing the rumor, and hence the progress of rumor propagation is stalled. \blacktriangleleft

- **Theorem 4.** *There is a graph $G = (V, E)$ of size $\Theta(n)$, such that for any $S_0 = \{s\} \subset V$:*
- *In expectation, the run-time of adversarial RPULL is in $\Omega(\sqrt{n})$.*
 - *W.h.p., the run-time of randomized RPULL is in $O(\log^2 n)$.*

Proof. Let G' and G'' be duplicates of the graph G from **Lemma 3**, r'_α and r''_α being the respective duplicates of r_α . We set $G := G' \cup G''$ and add the edge $\{r'_\alpha, r''_\alpha\}$. Without loss of generality let $s \in V'$.

In the random version, the rumor propagates through all of G' in $O(\log^2 n)$ rounds. Due to its low degree, r''_α gets the rumor from r'_α within $O(\log n)$ time after G' is informed and again, in $O(\log^2 n)$ rounds G'' is informed completely.

In the adversarial version, G'' can only learn the rumor from G' through edge $\{r'_\alpha, r''_\alpha\}$. But once r''_α knows the rumor, we can apply Lemma 3 again to prove that now progress is stalled. \blacktriangleleft

3 Comparison of PULL and RPULL

In this section we compare the two algorithms PULL and random RPULL on general graphs, i.e., we analyze how many rounds of random RPULL are enough to cover the progress of one round of PULL. More precisely, we show that w.h.p. the set of nodes informed after $O(\frac{\Delta}{\delta} \cdot \log n)$ rounds of random RPULL is a superset of the set of nodes informed in a single PULL round. We manage to do so by coupling both algorithms. At the end of the section we head out to prove that this bound is tight. Whenever we talk about RPULL in this section we mean random RPULL.

3.1 Dominance and Couplings

We begin with two examples of insufficient definitions of domination between two rumor spreading algorithms.

Showing for two algorithms \mathcal{A} and \mathcal{A}' that $\mathbf{P}(u \in S^{\mathcal{A}}) \geq \mathbf{P}(u \in S^{\mathcal{A}'})$ holds for all $u \in U$ is not enough to obtain a natural dominance definition of \mathcal{A} over \mathcal{A}' , since due to dependencies for a set M with $|M| > 1$ it might still be true that $\mathbf{P}(M \subseteq S^{\mathcal{A}}) < \mathbf{P}(M \subseteq S^{\mathcal{A}'})$.

Showing that $\mathbf{P}(M \subseteq S^{\mathcal{A}}) \geq \mathbf{P}(M \subseteq S^{\mathcal{A}'})$ (*) holds for all $M \subseteq U$ is not enough either. Assume the following example: Let $U = \{a, b, c\}$ be the set of uninformed nodes. Assume that under \mathcal{A} the probability that the set of newly informed nodes equals $\{a, b, c\}$, $\{a\}$, $\{b\}$ or $\{c\}$ is $1/8 + \epsilon$ each and the probability that it equals one of the sets $\{a, b\}$, $\{a, c\}$, $\{b, c\}$ or \emptyset is $1/8 - \epsilon$ each. Under \mathcal{A}' we inform any of those sets with probability $1/8$. A direct computation for all $M \subseteq \{a, b, c\}$, e.g., for $M = \{a\}$, $\mathbf{P}(\{a\} \subseteq S^{\mathcal{A}'}) = 1/2$ and

$$\mathbf{P}(\{a\} \subseteq S^{\mathcal{A}}) = \mathbf{P}(S^{\mathcal{A}} = \{a\}) + \mathbf{P}(S^{\mathcal{A}} = \{a, b\}) + \mathbf{P}(S^{\mathcal{A}} = \{a, c\}) + \mathbf{P}(S^{\mathcal{A}} = \{a, b, c\}) = 1/2,$$

shows that inequality (*) is fulfilled for any $M \subseteq U$, but the probability of the event “at least 2 nodes are informed” is by 2ϵ smaller for \mathcal{A} than for \mathcal{A}' . In the following we introduce the classical method to relate stochastic processes, i.e., the notion of (first order) stochastic dominance. However, we show that (proper) stochastic dominance between RPULL and PULL does not exist and thus we weaken the notion afterwards.

Stochastic Dominance and Coupling

Let $(\mathcal{S}, \preceq_{\mathcal{S}})$ be a finite distributive lattice and let X_1 and X_2 be random variables with distributions \mathbf{P}_1 and \mathbf{P}_2 which take values in \mathcal{S} . A function $f : \mathcal{S} \rightarrow \mathbb{R}$ is called *increasing* if $A \preceq_{\mathcal{S}} B$ implies $f(A) \leq f(B)$.

► **Definition 5** (Stochastic Dominance). We say that X_2 *stochastically dominates* X_1 if $E(f(X_2)) \geq E(f(X_1))$ holds for *every* increasing function $f : \mathcal{S} \rightarrow \mathbb{R}$, where $E(\cdot)$ denotes the expected value.

Alternative to Definition 5, one can show that one process stochastically dominates a second process by defining a monotone coupling between the processes (cf. Theorem 7).

► **Definition 6** ((Monotone) Coupling). A coupling of two random processes X_1 and X_2 , taking values in \mathcal{S} with distributions \mathbf{P}_1 and \mathbf{P}_2 , is a joint distribution $\hat{\mathbf{P}}$ of a random process (\hat{X}_1, \hat{X}_2) taking values in $\mathcal{S} \times \mathcal{S}$, such that its marginals equal the distributions of X_1 and X_2 , respectively, i.e.,

$$\begin{aligned} \sum_{B \in \mathcal{S}} \hat{\mathbf{P}} \left((\hat{X}_1, \hat{X}_2) = (A, B) \right) &= \mathbf{P}_1(X_1 = A) \quad \forall A \in \mathcal{S} \text{ and} \\ \sum_{A \in \mathcal{S}} \hat{\mathbf{P}} \left((\hat{X}_1, \hat{X}_2) = (A, B) \right) &= \mathbf{P}_2(X_2 = B) \quad \forall B \in \mathcal{S}. \end{aligned}$$

A coupling $\hat{\mathbf{P}}$ is called *monotone* (written $X_1 \leq X_2$) if additionally the following holds:

$$\forall A, B \in \mathcal{S} \text{ with } \hat{\mathbf{P}} \left((\hat{X}_1, \hat{X}_2) = (A, B) \right) > 0 \text{ it follows that } A \preceq_{\mathcal{S}} B. \quad (1)$$

Note that the choice of a coupling between two processes is generally not unique. The following theorem, Strassen's Theorem [19, 7], shows an equivalence between stochastic dominance and the notion of monotone couplings.

► **Theorem 7** (Strassen [19, 7]). *The following are equivalent:*

1. X_2 stochastically dominates X_1 ,
2. *There exists a monotone coupling between X_1 and X_2 such that $X_1 \leq X_2$,*
3. $\mathbf{P}(X_2 \in F) \geq \mathbf{P}(X_1 \in F)$ holds for every monotone set $F \subseteq \mathcal{S}$.³

Stochastic dominance/monotone couplings are the commonly known method to relate stochastic processes and we would like to show that $O(\frac{\Delta}{\delta} \log n)$ rounds of random RPULL stochastically dominate one round of PULL. This, however, is not possible as one can easily construct a graph in which some node u is informed with probability 1 in one round of PULL, but with probability less than 1 in $O(\frac{\Delta}{\delta} \log n)$ rounds of RPULL.⁴ Therefore we introduce the notion of highly probable monotone couplings and – in analogy to the equivalencies from Strassen's Theorem – also the notion of highly probable stochastic dominance.

► **Definition 8.** A coupling $\hat{\mathbf{P}}$ of two random processes X_1 and X_2 is called **monotone w.h.p.** (w.r.t. to n) (written $X_1 \leq_{\text{w.h.p.}} X_2$) if for some $c > 1$ it satisfies

$$\sum_{A \not\preceq B} \hat{\mathbf{P}} \left((\hat{X}_1, \hat{X}_2) = (A, B) \right) \leq \frac{1}{n^c}. \quad (2)$$

We say X_2 *stochastically dominates X_1 with high probability* (w.r.t. n), if there exists a coupling between X_1 and X_2 that is monotone with high probability ($X_1 \leq_{\text{w.h.p.}} X_2$).

³ A set $F \subseteq \mathcal{S}$ is called monotone if $A \in F$ and $A \preceq_{\mathcal{S}} B$ implies $B \in F$.

⁴ Figure 2 in Appendix A can be used to verify this.

In this paper we will set $\mathcal{S} = 2^U$ to be the power set of U , where $U \subseteq V$ is the set of uninformed nodes, $\preceq_{\mathcal{S}}$ equals the subset relation on U and X_2 and X_1 will be the respective random variables describing which nodes get informed in RPULL and PULL. In general, the parameter n in Definition 8 can be freely chosen; in our setting n will be the number of nodes of the communication network. With this choice of parameters a monotone coupling (w.h.p. w.r.t. n) is the desired relation of PULL and RPULL.

3.2 W.h.p. Monotone Coupling between PULL and RPULL

► **Theorem 9.** *W.h.p., for any set of informed nodes $S \subseteq V$, $T = O\left(\frac{\Delta}{\delta} \log n\right)$ rounds of random RPULL stochastically dominate a single round of PULL.*

► **Corollary 10.** *If in a graph G with initially informed nodes $S \subseteq V$ the PULL algorithm informs all nodes in \mathcal{T} rounds with probability p , then the random RPULL algorithm informs all nodes in time $O\left(\mathcal{T} \cdot \frac{\Delta}{\delta} \cdot \log n\right)$ with probability $(1 - o(1))p$.*

By PUSH – RPULL we denote the combination of RPULL with a simultaneous execution of the classic PUSH protocol. The restriction of a single node to answer only a limited number of requests does not limit the progress of the PUSH algorithm when disseminating a rumor. Hence we deduce the following corollary.

► **Corollary 11.** *W.h.p., for any set of informed nodes $S \subseteq V$, $T = O\left(\frac{\Delta}{\delta} \log n\right)$ rounds of PUSH – RPULL stochastically dominate a single round of PUSH – PULL.*

To reduce dependencies between nodes which request from the same neighbor we introduce a new algorithm VPULL (virtual pull). Note that VPULL is only introduced as a tool to analyze the process RPULL; difficulties/impossibilities that arise in an actual implementation of VPULL are not relevant. The proof of Theorem 9 is then split into two parts:

1. Lemma 13: W.h.p., T rounds of RPULL stoch. dominate one execution of VPULL,
2. Lemma 16: One execution of VPULL stochastically dominates one round of PULL.

Then Theorem 9 follows from the transitivity of the stochastic dominance relation.

By RPULL_T we denote the (randomized) process RPULL which runs for T rounds, by VPULL we denote one execution of VPULL and by PULL_1 we denote the process PULL which runs for one round only. The random variables S_T^{RPULL} , S^{VPULL} and S_1^{PULL} denote the respective sets of nodes that are informed after the corresponding number of rounds. The processes RPULL_T , VPULL and PULL_1 are not completely characterized by the random variables S_T^{RPULL} , S^{VPULL} and S_1^{PULL} – one has to include information about all requests and messages, that are sent by all nodes, to fully describe the random processes. Nevertheless, to show the desired result, it is sufficient to find a monotone coupling where condition (1) and (2), respectively, are fulfilled with regard to the subset relation of the set valued random variables S_T^{RPULL} , S^{VPULL} and S_1^{PULL} .

Definition of VPULL

Let us define *weakly connected* nodes $u \in U$ as nodes for which $d_S(u)/d(u) \leq 1/2$ and *strongly connected* otherwise. To introduce the process VPULL we need the following parameters, which we fix later.

$$K = \Theta\left(\frac{\Delta}{\delta} + \log n\right), \quad T' = O\left(\frac{\Delta}{\delta} \log n\right) \quad \text{and} \quad T = \Theta(T'), \quad T \gg T'$$

An execution of VPULL consists of two phases. In the first phase nodes send tokens instead of the actual rumor and w.h.p. nodes who have received a token in the first phase are informed at the end of the second phase. In an execution of VPULL we let $X_v(t)$ be the number of tokens which node v has sent up to round t . In a specific round t denote with $R_v(t)$ the set of nodes requesting from some informed node $v \in S$ and with $r_v(t) = |R_v(t)|$ its cardinality. R_v , r_v and X_v are random variables which describe certain properties of an execution of VPULL, where large values of X_v or r_v indicate the unlikely case in which the (strict) monotonicity of a tentative coupling between VPULL and RPULL $_T$ might break.

► **Definition 12** (Good, Bad Execution). An execution of VPULL is called a *bad execution* if for some $v \in V$ or $1 \leq t \leq T$ it holds that $X_v(t) > K$ or $r_v(t) > K$, otherwise it is called a *good execution*.

A formal definition of VPULL can be found in Appendix, Section B. Here, we describe the algorithm informally. An execution of VPULL is split into two phases – the first phase consists of T rounds and the second phase of one round. In the first phase an uninformed node requests the rumor uniformly at random from one of its neighbors and an informed node v decides with probability $\frac{r_v}{T}$ whether to send out a token – in which case it selects, uniformly at random, one of its incoming requests as destination for the token. Nodes that get a token in those T rounds, stop requesting from neighbors, but are still unable to forward any information to neighbors in consecutive rounds. In the second phase the limit to the number of requests that can be served by an informed node is stripped away. Then, in case of a bad execution (note that the first phase determines whether an execution is good or bad) all actions from the first T rounds are discarded and all uninformed nodes perform one round of PULL. In case of a good execution all uninformed strongly connected nodes perform one round of PULL and afterwards all nodes holding a token are being informed. If we assume that tokens are as valuable as the information itself, in each of the T rounds of the first phase, VPULL differs from RPULL only in the fact that the selected incoming connection is established with probability $\frac{r_v}{T}$ whereas it is established deterministically in RPULL. For an uninformed node $u \in U$, that chooses to request a neighbor $v \in S$, this normalizes the probability to get a token to $1/T'$, independent of the amount of other requesting nodes. Except for the second phase this algorithm is clearly dominated by RPULL.

W.h.p. Monotone Coupling between RPULL and VPULL

First, we generate a monotone coupling between T rounds of RPULL and the first phase of VPULL as follows: For each round in the first phase both processes use the same randomness to decide on the outgoing calls of uninformed nodes (if in a round $t > 1$ a node is uninformed in VPULL but not in RPULL the process VPULL uses additional randomness; the contrary cannot happen). VPULL uses additional randomness to decide whether a node, which is contacted, sends out any message at all, confer line 13 from Algorithm 1. This simultaneous execution of both algorithms gives rise to a coupling of the first phase of VPULL and RPULL $_T$. Clearly, a node that is provided with a token in VPULL in any round is then also informed in RPULL, i.e., the coupling is monotone.

► **Lemma 13.** RPULL $_T$ *stochastically dominates* VPULL *with high probability*.

Proof. Under the assumption that tokens are as valuable as the information itself we constructed a monotone coupling of the first phase of VPULL and S_T^{RPULL} . Then, it is sufficient to prove that in the second phase of VPULL, w.h.p., no node is informed, that has not been informed in the T rounds of RPULL: If neither ever any value r_v nor any X_v

exceeded K , then only strongly connected nodes simulate one round of PULL in the second phase of the VPULL algorithm. We claim that, w.h.p., each strongly connected node has been informed in the first T rounds of RPULL.

A strongly connected node $u \in U$ requests from an informed node $v \in S$ with probability at least $1/2$. In any given round due to Markov inequality with probability at least $1/2$ no more than $2\Delta/\delta$ nodes $u' \in U$ connect to v . The probability for u to get informed under RPULL is thus at least $\frac{\delta}{8\Delta}$. Choosing $T = O\left(\frac{\Delta}{\delta} \log n\right)$ big enough and a union bound gives us that, w.h.p., all strongly connected nodes are informed in process RPULL_T .

To conclude, we prove that w.h.p. neither r_v nor X_v exceed K for any node v during an execution of VPULL. Let $1 < \kappa < c_{T,T'}$ be constants, $K' := \frac{\Delta}{\delta} + \log n$ and $K = c_{T,T'}K'$.

W.h.p., $r_v \leq K$ in VPULL_T for all v . For a fixed informed node v , in expectation, no more than $\frac{\Delta}{\delta}$ nodes can request from v . Using a Chernoff bound for a single round and a single node, $\mathbf{P}\left(r_v \geq \kappa\left(\frac{\Delta}{\delta} + \log n\right)\right) \leq n^{-\Theta(\kappa)}$ holds. With a union bound over all nodes and all rounds and κ large enough we obtain that, w.h.p., r_v never exceeds $\kappa K'$ and therefore neither K . A union bound over all nodes concludes the proof.

W.h.p., $X_v \leq K$ in VPULL_T for all v . For a fixed v , note that, w.h.p., in a single round no more than $\kappa K'$ nodes request from v , and therefore, X_v is increased at most with probability $\kappa K'/T'$ in any round. Over T rounds, in expectation, no more than $\kappa K' \frac{T}{T'}$ increments of X_v happen, and again a Chernoff bound gives us that X_v does not exceed $2\kappa K' \frac{T}{T'}$ with high probability. Choosing $c_{T,T'} = 2\kappa \frac{T}{T'}$ and a union bound over all nodes concludes the proof. \blacktriangleleft

Stochastic Dominance between VPULL and PULL

In a single round of PULL a node $u \in U$ is informed with probability $\frac{d_S(u)}{d(u)}$, independently from which other nodes are informed. For one execution of VPULL we can show that a node u is also informed at least with probability $\frac{d_S(u)}{d(u)}$, independently from which other nodes get informed (Lemma 14). Afterwards, we prove that Lemma 14 is sufficient to deduce the stochastic dominance of one execution of VPULL over PULL_1 . For $u \in U$ and random process X , let \mathcal{C}_u^X be the set of all conditions of the type $v \in X$ or $v \notin X$ where $v \neq u$.

► **Lemma 14.** *In VPULL a node $u \in U$ is informed at least with probability $\frac{d_S(u)}{d(u)}$, independently from which other nodes are informed, i.e., for all sets of conditions $I \subseteq \mathcal{C}_u^{\text{VPULL}}$ and $J \subseteq \mathcal{C}_u^{\text{PULL}}$ with $\mathbf{P}(I), \mathbf{P}(J) > 0$ the following holds*

$$\mathbf{P}\left(u \in S_{T+1}^{\text{VPULL}} \mid I\right) \geq \frac{d_S(u)}{d(u)} = \mathbf{P}\left(u \in S_1^{\text{PULL}}\right) = \mathbf{P}\left(u \in S_1^{\text{PULL}} \mid J\right). \quad (3)$$

Proof. If $u \in U$ is strongly connected, the result holds because VPULL executes one round of PULL for u in either way. In a bad execution, VPULL executes one round of PULL for any uninformed node and the claim holds trivially. Thus assume that u is weakly connected and we are in a good execution. Let $s = d_S(u)$ and $N_S(u) = \{v_1, \dots, v_s\}$ be the neighbors of u in S . We call a node $v \in N_S(u)$ *busy* w.r.t. u in round t if it informs some node other than u . Let y_t be the number of busy nodes in round t w.r.t. u . In a good execution (which we denote by \mathbb{G}), any node in $N_S(u)$ can inform at most K nodes and hence there is the following constraint on the sum of all y_t 's

$$\sum_{t=1}^T y_t \leq s \cdot K. \quad (4)$$

We can ignore conditions in I corresponding to nodes which do not have a common neighbor with $N_S(u) \cup \{u\}$ because u can only get the rumor directly through S . The only negative effect on the probability that u gets informed by the conditions in I can be captured by the number of busy nodes w.r.t. u . However, since the number of nodes which are informed per node in a good execution is small compared with T , there are sufficiently many rounds with sufficiently many non-busy nodes to inform u . More precisely, if u requests from a non-busy node it is informed at least with probability $\frac{1}{T'}$. Thus, the probability that u , conditioned on $I \wedge \mathbb{G}$ with $\mathbf{P}(I \wedge \mathbb{G}) > 0$, is not informed is smaller or equal to (with $c = T/T'$)

$$\prod_{t=1}^T \left(1 - \frac{s - y_t}{d(u) \cdot T'}\right) \leq \left(1 - \frac{s(1 - \frac{K}{T})}{d(u) \cdot T'}\right)^T \leq e^{-c(1 - \frac{K}{T}) \frac{s}{d(u)}} \leq 1 - \frac{d_S(u)}{d(u)}.$$

The first inequality holds because under constraint (4) the expression on the left hand side is maximized for $y_t = \frac{s \cdot K}{T}$. The last inequality holds due to $\frac{s}{d(u)} \leq 1/2$, $c(1 - K/T) \geq 2$ and the fact that $e^{-2x} \leq 1 - x$ for any $x \in [0, 1/2]$. \blacktriangleleft

Lemma 14 is sufficient to show stochastic domination of one execution of VPULL over PULL₁. For a formal proof one can either directly construct a monotone coupling between the random variables or use the following result due to Holley [15].

► **Theorem 15** (Holley Inequality, [15]). *Let $(\mathcal{S}, <)$ be a distributive lattice and let μ_1, μ_2 be measures on this lattice. The Holley criterion is satisfied if*

$$\mu_1(A \cap B) \mu_2(A \cup B) \geq \mu_1(A) \mu_2(B) \text{ holds for all } A, B \in \mathcal{S}. \quad (5)$$

If the Holley criterion is satisfied for μ_1 and μ_2 then

$$\sum_{A \in \mathcal{S}} \mu_1(A) f(A) \geq \sum_{A \in \mathcal{S}} \mu_2(A) f(A) \text{ holds for all increasing functions } f : \mathcal{S} \rightarrow \mathbb{R}. \quad (6)$$

► **Lemma 16.** *One execution of VPULL stochastically dominates PULL₁.*

Proof. For the proof let U be those uninformed nodes u with $0 < d_S(u) < d(u)$ and consider the distributive lattice $(\mathcal{S}, \preceq_{\mathcal{S}}) = (2^U, \subseteq)$. Every uninformed node u which is not contained in this redefined U has either no connection to S at all, i.e., it is not informed in either process, or $d_S(u) = d(u)$ holds, i.e., it is informed with probability one in either process because also VPULL executes one round of PULL for it. Hence it is sufficient to show stochastic domination of S^{VPULL} over S_1^{PULL} restricted to this redefined set U . This choice of U provides $0 < \mathbf{P}(S_1^{\text{PULL}} = A), \mathbf{P}(S^{\text{VPULL}} = A) < 1$ for all $A \in \mathcal{S}$ and we define the *strictly positive* measures $\mu_1(F) := \mathbf{P}(S^{\text{VPULL}} \in F)$ and $\mu_2(F) := \mathbf{P}(S_1^{\text{PULL}} \in F)$ for $F \subseteq 2^U$. For A in $\mathcal{S} = 2^U$, $x \in U$ define $A^x := A \cup \{x\}$ and $A_x := A \setminus \{x\}$.

The proof of Lemma 16 then follows with Theorem 15, the following claims and the definition of the expected value of an increasing function $f : \mathcal{S} \rightarrow \mathbb{R}$.

► **Claim 17** (Quotient Rule).

$$\frac{\mu_1(A^x)}{\mu_1(A_x)} \geq \frac{\mu_2(B^x)}{\mu_2(B_x)} \text{ holds for all } A, B \in \mathcal{S}. \quad (7)$$

► **Claim 18.** *The quotient rule (17) implies that the Holley criterion is satisfied for the measures μ_1 and μ_2 .*

The proofs of both claims are based on Lemma 14 and can be found in the appendix. \blacktriangleleft

Proof of Theorem 9. The proof is a direct combination of Lemma 13 and Lemma 16. \blacktriangleleft

3.3 The Round-by-Round Analysis is Tight

► **Lemma 19.** *The time bound $T = O(\frac{\Delta}{\delta} \log n)$ from Theorem 9 is tight.*

Proof. In order for T rounds of RPULL to dominate one round of PULL (w.h.p), any node v must get informed in T rounds of RPULL with at least the same probability as within one round of PULL (or w.h.p., if the latter probability equals one). We construct a graph G for which at least $T = \Omega(\frac{\Delta}{\delta} \log n)$ rounds of RPULL are necessary to guarantee this. A picture depicting the graph can be found in Appendix A.

We partition set V into $V = A \cup B \cup T_{1,1} \cup \dots \cup T_{k^2, k^2}$, where $k = n^{1/5}$, $A = \{a_1, \dots, a_{k^2}\}$ and $B = \{b_1, \dots, b_{k^2}\}$. A and B form a complete bipartite graph with edges running between A and B . For each $j \in [k^2]$, node b_i is connected to *one* node $t_{i,j} \in T_{i,j}$. Each $T_{i,j}$ forms a complete graph of size k . In this graph, $\delta = k - 1$ (acquired in $T_{i,j}$) and $\Delta = 2k^2$ (nodes in B), and therefore $\Delta/\delta \in \Theta(k)$. The total size of the graph is $|V| = n + o(n)$. Initially, we let $S_0 = B$.

In this graph, within one round of PULL, all nodes of A are informed with probability 1. Now, consider the same graph after $m \leq k^2/2$ rounds of RPULL and let us assume that some node $a \in A$ is still uninformed. It requests in this round from some node b_i . Let X_i be the number of requests at b_i . Within m rounds, each node b_i managed to inform at most m of its neighbors from $NB_i := \{t_{i,1}, \dots, t_{i,m}\}$. Since $m \leq k^2/2$, at least half of all nodes in NB_i are still uninformed and thus, since they have degree k , $\mathbf{E}[X_i] \geq k/2$. Applying Chernoff, we get that w.h.p., $X_i \geq k/4$. In this scenario for a the probability to be chosen over one of its competitors is at most $4/k$, regardless of m , and therefore, $\mathbf{P}(a \in S_m^{\text{RPULL}}) \leq (1 - \frac{4}{k})^m$. For this to fall below $1/n$, m has to be in $\Theta(\frac{\Delta}{\delta} \log n)$. ◀

4 Conclusions

Lemma 19 and Theorem 9 show that to simulate one round of PULL, $\Theta(\frac{\Delta}{\delta} \log n)$ rounds of RPULL are required. However, in case one wants stochastic dominance (w.h.p.) over $\mathcal{T} > 1$ rounds of PULL, the lower bound proof of Lemma 19 does not hold anymore. We believe that for $\mathcal{T} = \Omega(\log n)$, on any graph G and any set of initially informed nodes $S \subseteq V$, $O(\mathcal{T}(\frac{\Delta}{\delta} + \log n))$ or maybe even $O(\mathcal{T}(\frac{\Delta}{\delta}) + \log n)$ rounds of RPULL suffice to stochastically dominate \mathcal{T} rounds of PULL. That proving this assumption might be a challenging task is underlined by a similar conjecture in [1], in which the authors do a coupling of synchronous and asynchronous PUSH-PULL. They obtain a similar multiplicative $O(\log n)$ factor and also conjecture that it can be improved to an additive $O(\log n)$ term.

A possible alternative restriction of the PUSH-PULL protocol could be given by the following algorithm. In each round, every node requests from an outgoing neighbor chosen uniformly at random. At each node, one of the incoming requests is chosen (e.g., uniformly at random) and a connection to the requesting node is established. Finally, over all established links between an informed and an uninformed node, the uninformed node learns the rumor. Note that unlike in the restricted PUSH-PULL variant described in our paper, here, also two informed nodes or two uninformed nodes could be paired. Such a PUSH-PULL variant can be analyzed in an analogous way to our analysis of the RPULL protocol and it can be shown that $O(\frac{\Delta}{\delta} \log n)$ rounds of this algorithm stochastically dominate a single round of the regular PUSH-PULL protocol.


References

- 1 Hüseyin Acan, Andrea Collecchio, Abbas Mehrabian, and Nick Wormald. On the push&pull protocol for rumour spreading. *CoRR*, abs/1411.0948, 2014.
- 2 K. Censor-Hillel, B. Haeupler, J. A. Kelner, and P. Maymounkov. Global computation in a poorly connected world: Fast rumor spreading with no dependence on conductance. In *Proc. Symp. on Theory of Computing (STOC)*, pages 961–970, 2012.
- 3 K. Censor-Hillel and H. Shachnai. Fast information spreading in graphs with large weak conductance. *SIAM J. Computing*, 41(6):1451–1465, 2012.
- 4 J. Chen and G. Pandurangan. Optimal gossip-based aggregate computation. In *Symp. on Parallelism in Algorithms and Architectures (SPAA)*, pages 124–133, 2010.
- 5 Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proc. Symp. on Theory of Computing (STOC)*, pages 399–408, 2010.
- 6 A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database management. In *Proc. Symp. on Principles of Dist. Computing (PODC)*, pages 1–12, 1987.
- 7 C. Domb and J.L. Lebowitz, editors. *Phase transitions and critical phenomena, Volume 18*. Academic Press, 2001.
- 8 U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.
- 9 N. Fountoulakis, K. Panagiotou, and T. Sauerwald. Ultra-fast rumor spreading in social networks. In *Proc. Symp. on Discrete Algorithms (SODA)*, pages 1642–1660, 2012.
- 10 A. M. Frieze and G. R. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Appl. Math.*, 10(1):57–77, 1985.
- 11 G. Giakkoupis and T. Sauerwald. Rumor spreading and vertex expansion. In *Proc. Symp. on Discrete Algorithms (SODA)*, pages 1623–1641, 2012.
- 12 George Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *Proc. Int. Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 57–68, 2011.
- 13 George Giakkoupis. Tight bounds for rumor spreading with vertex expansion. In *Proc. Symp. on Discrete Algorithms (SODA)*, pages 801–815, 2014.
- 14 G.R. Grimmett. *The Random-Cluster Model*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2006.
- 15 Richard Holley. Remarks on the fkg inequalities. *Communications in Mathematical Physics*, 36(3):227–231, 1974.
- 16 R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proc. Symp. on Foundations of Computer Science (FOCS)*, pages 565–574, 2000.
- 17 D. Kempe, A. Dorba, and J. Gehrke. Gossip-based computation of aggregate information. In *Proc. Symp. on Foundations of Computer Science (FOCS)*, pages 482–491, 2003.
- 18 D. Mosk-Aoyama and D. Shah. Fast distributed algorithms for computing separable functions. *IEEE Trans. on Inf. Theory*, 54(7):2997–3007, 2008.
- 19 Volker Strassen. The Existence of Probability Measures with Given Marginals. *Annals of Mathematical Statistics*, 36(2):423–439, 1965.

Appendix

A Pictures

A.1 Picture for Lemma 3, Figure 1




■ **Figure 1** Picture for Lemma 3, proving that random and adversarial RPULL have exponentially different running times on general graphs. Grey areas indicate fully connected parts of the graph.

In random RPULL the node r_ζ learns the rumor within $O(\log^2 n)$ rounds and can spread the information through the graph in polylogarithmic time.

In adversarial RPULL the adversary prevents r_ζ from learning the rumor by always disseminating the rumor to one of the requesting nodes of in D_1, \dots, D_m in every round. We can show that the number of informed D_i s grows slowly and hence such requests exist w.h.p. as long as no node in D_ζ is informed. Also, with only few D_i s informed, due to their high degrees, leaf nodes in L_ζ are unlikely to request from a D_i containing the rumor, and hence the progress of rumor propagation is stalled.

A.2 Picture for Lemma 19, Figure 2



■ **Figure 2** Picture for Lemma 19, proving that $\Omega(\frac{\Delta}{\delta} \log n)$ rounds of RPULL are necessary to simulate one round of PULL.

All not filled circular nodes in the bipartite graph, i.e., the nodes in B , have the rumor at the start of the execution. In one round of PULL all filled nodes in the bipartite graph (A) learn the rumor with probability one. In random RPULL $\Omega(\frac{\Delta}{\delta} \log n)$ rounds are necessary to inform these nodes, because for each of the nodes in A the probability to be informed in one round of RPULL is in $\Theta(1/k)$. This is due to the high number of requests from nodes in $T_{i,j}$ each round to the informed nodes in B .

B Pseudocodes

Algorithm 1 One execution of VPULL ($(T + 1)$ -rounds)

Input: K – threshold for bad execution; T' – parameter to normalize probabilities

States: informed; uninformed

Oracle knowledge: $d_S(v)$ for every node v ; $BE := \bigvee_{v \in V} BE_v$

Variables: R_v set of nodes requesting from v in the corresp. round ($r_v := |R_v|$)
 BE_v boolean indicator for bad execution caused at node v
 $tokenReceived_v$ indicates whether a node will be informed after T rounds

```

1:  $BE_v \leftarrow \mathbf{false}$ ;  $tokenReceived_v \leftarrow \mathbf{false}$ 
2: for  $T$  rounds do
3:   switch  $state_v$  do
4:     case uninformed
5:       if  $tokenReceived_v = \mathbf{false}$  then
6:         send request for rumor uniformly at random
7:         if  $msg = token$  then
8:            $tokenReceived_v \leftarrow \mathbf{true}$ 
9:       case informed
10:      if  $r_v > K$  or  $X_v > K$  then           // bad execution has been detected locally
11:         $BE_v \leftarrow \mathbf{true}$ 
12:      else
13:        with probability  $r_v/T'$  do
14:          send  $token$  to uniformly at random chosen node in  $R_v \neq \emptyset$ 
15:           $X_v \leftarrow X_v + 1$ 
        // Round  $T + 1$ :
16: request  $(BE, d_S(v))$  from global oracle
17: if  $BE = \mathbf{true}$  then           // bad execution has been detected globally
18:   execute one round of PULL           // executed locally
19: else
20:   if  $d_S(v)/d(v) > 1/2$  then           // node is strongly connected
21:     execute one round of PULL           // executed locally
22:   else if  $tokenReceived_v = \mathbf{true}$  then
23:      $state_v \leftarrow \mathbf{informed}$            // node learns rumor

```

Note that VPULL is only introduced as a tool and we do not claim that it can be implemented as a distributed algorithm. Correct values for the parameters K and T' are defined in section 3.2. The variables X_v , R_v , r_v , BE_v and BE can either be understood as random variables describing an execution of the VPULL algorithm or they can be updated directly in the algorithm as done above. Except for line 16 which uses global knowledge VPULL can be seen as a distributed algorithm.

C Proofs

Chernoff Statement for Lemma 2

► **Lemma 20.** Let X_1, X_2, \dots, X_n be independent geometric random variables with $X_i \sim \text{Geo}(p_i)$ for $i = 1, \dots, n$, $t \in \mathbb{N}$ and $0 < p_1 \leq p_2 \leq \dots < 1$. Let $X = \sum_{i=1}^n X_i$ and $\mu := \mathbf{E}[X]$. Then

$$\mathbf{P}(X > 3(\mu + t)) \leq e^{-\frac{p_1}{2}\mu - p_1 t}. \quad (8)$$

Proof. Let $\gamma = -\ln(1 - \frac{p_1}{3}) > 0$. Because of $p_1 \leq p_i$ this implies $e^\gamma \cdot (1 - p_i) < 1$ for all $i = 1, \dots, n$. We need this condition at (*) in the proof of the following claim.

► **Claim 21.** $\mathbf{E}[e^{\gamma X_i}] \leq 1 + \frac{p_1}{2p_i}$.

Proof of Claim 21. With a straight forward calculation one obtains

$$\begin{aligned} \mathbf{E}[e^{\gamma X_i}] &= \sum_{k=1}^{\infty} \mathbf{P}(X_i = k) e^{\gamma k} = p_i e^\gamma \sum_{k=1}^{\infty} ((1 - p_i) e^\gamma)^{k-1} \\ &\stackrel{(*)}{=} \frac{p_i e^\gamma}{1 - (1 - p_i) e^\gamma} = \frac{p_i}{e^{-\gamma} - 1 + p_i} \\ &= 1 + \frac{1 - e^{-\gamma}}{e^{-\gamma} + p_i - 1} = 1 + \frac{1 - (1 - \frac{p_1}{3})}{1 - \frac{p_1}{3} + p_i - 1} \\ &= 1 + \frac{\frac{p_1}{3}}{p_i - \frac{p_1}{3}} \leq 1 + \frac{p_1}{2p_i}. \end{aligned}$$

◀

Now we can show the actual result. Since $x \mapsto e^{\gamma x}$ is an increasing function for $\gamma > 0$ we obtain

$$\begin{aligned} \mathbf{P}(X > 3(\mu + t)) &= \mathbf{P}(e^{\gamma X} > e^{\gamma 3(\mu + t)}) && \text{(Markov)} \\ &\leq e^{-3\gamma(\mu + t)} \prod_{i=1}^n \mathbf{E}[e^{\gamma X_i}] && \text{(claim 1)} \\ &\leq e^{-3\gamma(\mu + t)} \prod_{i=1}^n \left(1 + \frac{p_1}{2p_i}\right) && (1 + x \leq e^x, x \in \mathbb{R}) \\ &\leq \left(1 - \frac{p_1}{3}\right)^{3(\mu + t)} e^{\sum_{i=1}^n \frac{p_1}{2p_i}} && \left(\mu = \sum_{i=1}^n \frac{1}{p_i}\right) \\ &= \left(1 - \frac{p_1}{3}\right)^{3(\mu + t)} (e^{\frac{p_1}{2}\mu}) && (1 - x \leq e^{-x}, x \in \mathbb{R}) \\ &\leq e^{-(p_1(\mu + t) + \frac{p_1}{2}\mu)} = e^{-\frac{p_1}{2}\mu - p_1 t}. \end{aligned}$$

◀

C.1 Proofs for Section 2

► **Lemma 2.** Let G be a tree network with $S_0 = \{r\}$. Then in both random and adversarial RPULL it takes $O(\max_{\text{path}_p} M_p + \Delta \log n)$ rounds to fully inform all nodes in V , w.h.p..

Proof. The analysis we use does hold for adversarial RPULL.

First look at a path $p = p_{r,l} = (r = v_0, v_1, v_2, \dots, v_q = l)$ from the root to some leaf l and define $\Delta_p = \max_{v \in p} d(v)$. Let T_i be the random variable that indicates the round in which node v_i gets informed and with $T_0 := 0$ we can define $X_i := T_i - T_{i-1}$ for $m = 1, \dots, q$, the time node v_{i-1} needs to pass the information forward to node v_i . For simplicity in the following we define $u := v_i$ and $s := v_{i-1}$.

Once s gets informed, a round is called free if no node in $N(s) \setminus \{v_{i-2}, u\}$ requests to s , otherwise it is called congested. In a free round, u gets the rumor with probability at least $1/d(u)$, i.e., the number of free rounds is upper bounded by a geometric random variable $Y_{i,\text{free}} \sim \text{Geom}(1/d(u))$. On the other hand, in a congested round, at least one neighbor of s does get informed, so there can not be more than $d(s)$ such rounds.

In total we get for round t' that $\mathbf{P}(X_i \geq t') \leq \mathbf{P}(Y_{i,\text{free}} + d(s) \geq t')$ for node $u = v_i$ and with $T_q = \sum_{i=1}^q X_i$ we conclude that $\mathbf{P}(T_q \geq t') \leq \mathbf{P}(\sum_{i=1}^q Y_{i,\text{free}} + \sum_{i=1}^q d(v_{i-1}) \geq t')$.

Since we are on a tree, those geometric random variables are all independent, and we can apply the Chernoff Lemma 20. Let Δ_p be the largest degree of all nodes on the path excluding r , i.e., $p_1 = 1/\Delta_p$ in terms of the notation from Lemma 20. With $Y = \sum_{i=1}^q Y_{i,\text{free}}$ we have $\mu = D_p - d(r) \geq \Delta_p$ and we set $t = c\Delta_p \log n$ for some $c > 0$.

$$\mathbf{P}(Y > 3(\mu + t)) \leq e^{-\frac{\mu}{2\Delta_p}} e^{-\frac{c\Delta_p \log n}{\Delta_p}} < n^{-c},$$

i.e., w.h.p., v_q is informed within $O(D_p + \Delta_p \log n) = O(\max_{\text{path } p} D_p + \Delta \log n)$ rounds. In a tree there are at most n root-leaf paths, therefore a union bound over all individual paths concludes the proof. \blacktriangleleft

- **Lemma 3.** There is a graph $G = (V, E)$ of size $\Theta(n)$ with node $r_\alpha \in V$, $d(r_\alpha) \leq 3$, s.t.:
- For $S_0 = \{r_\alpha\}$, w.c.p., the run-time of adversarial RPULL is in $\Omega(\sqrt{n})$.
 - For any non-empty $S_0 \subset V$, w.h.p., the run-time of randomized RPULL is in $O(\log^2 n)$.

Proof. Recall that a picture depicting the graph can be found in Appendix A.

We use the notation w.h.p. (k) to say that some event holds with high probability with respect to k , i.e., with probability at least $1 - 1/k$.

We first introduce a graph type, with size depending on some parameter k , that we call a k -leaf-connected tree (k -LCT). In simple words, a k -LCT is a binary tree with k leaves, but with those k leaves being fully interconnected, i.e., forming a clique. Mathematically more precise, a graph $H = (B \cup L, E_B \cup E_L)$ with $|L| = k$ and $|B| = k - 1$ is a k -LCT, iff $H_L = (L, E_L)$ is a complete graph over k nodes and $H_B = (B \cup L, E_B)$ is a complete binary tree with its leaves being the nodes in L . While H is not strictly speaking a tree, we call nodes in L its leaves, $L(H) = L$ the *leaf set*, $B(H) = B$ its *branch set* and the root of H_B we call the *root* of H . Also, every node in H except for its root has a clearly designated *parent* (defined by H_B) and each node in B has two clearly designated *children* (with respect to the root in H_B).

- **Claim 22.** Let $G = (V, E)$ be a graph and $H = (V_H, E_H)$ be a subgraph of G that is a k -LCT. Furthermore, let any node $v \in V_H$ have at most one connection outside of V_H , i.e., $d_G(v) \leq d_H(v) + 1$. Then, w.h.p. (k) , for any non-empty set of nodes in H knowing the rumor, randomized RPULL informs all nodes in H within $O(\log k)$ rounds.

Proof of Claim 22. Without loss of generality let there be one node s having the rumor. If $s \in B := B(H)$, i.e., $d_G(s) \leq d_H(s) + 1 \leq 4$, then we can apply **Lemma 2** to get that *all* nodes in B are informed within $O(\log k)$ rounds. Let this be the case. All nodes in $L := L(H)$

have degree at most $k + 1$: $k - 1$ neighbors in L , one “parent-node” in B and at most one neighbor in $V \setminus V_H$. Each of them requests to its neighboring parent from B with probability at least $1/(k + 1)$, i.e., in each round, with probability at most $(1 - \frac{1}{k+1})^k < 1/2$, no node in L learns the rumor. By Chernoff, w.h.p.(k), after $O(\log k)$ rounds, at least one node in L knows the rumor. If $x > 0$ nodes in L are informed, then each uninformed node u in L requests from one of those x nodes (or a node in B) with probability at least $x/(k + 1)$ and with probability at least $(1 - 1/k)^{k-1} > 1/3$ node u is the only node requesting from its target. As long as $x < k/2$, with linearity of expectation, each round the expected number of newly informed nodes in L is in $\Omega(x)$. Once $x \geq k/2$ we can use a similar argument to show that, w.c.p., the number of uninformed nodes goes down by a constant factor each round. Hence, after $O(\log k)$ rounds in expectation, but also w.h.p.(k), all nodes in L are informed.

If initially $s \in L$, then with probability $p \in \{1/4, 1/3\}$ its parent node in B requests from s , while at the same time with probability at least $(1 - 1/k)^{k-1} > 1/3$ no other node in L requests from s . Hence, after $O(\log k)$ rounds, w.h.p., the parent node gets the rumor from s . The rest follows from reduction to the first case. \blacktriangleleft

We construct $G = (V_1 \cup V_\zeta, E)$ as follows. We let D_α and D_ζ be two n -LCTs, and we have m l -LCTs that we denote with D_1, D_2, \dots, D_m , where $l := \sqrt{n}$ and $m := c\sqrt{n}$ for some natural number c . We use the notation D_i for the corresponding k -LCT and its node set interchangeably. Their corresponding roots and leaf sets are denoted as $r_\alpha, r_\zeta, r_1, r_2, \dots, r_m$ and $L_\alpha, L_\zeta, L_1, L_2, \dots, L_m$ respectively, and with $l_{X,1}, l_{X,2}, \dots$ we enumerate the leaves of leaf set L_X . Let $C_\alpha = \{c_1, \dots, c_m\}$ be an arbitrary m -sized subset of D_α 's branch set B_α – for simplicity and in accordance to Figure 1 think of C_α as the layer of nodes in B_α that are at depth $\log m$.

We let $V_\zeta = D_\zeta$ and $V_1 = \{r\} \cup D_\alpha \cup D_1 \cup \dots \cup D_m$ and we add the following edges.

- Between r and D_ζ : We add an edge from r to r_ζ .
- Between r and D_α : For each $j \in [m \log n]$ we add an edge from r to $l_{\alpha,j}$.
- Between r and D_1, \dots, D_m : For each $i \in [m], j \in [\log n]$ we add an edge from r to $l_{i,j}$.
- Between D_1, \dots, D_m and D_ζ : For each $i \in [m]$ we add an edge from r_i to $l_{\zeta,i}$.
- Between D_1, \dots, D_m and C_α : For each $i \in [m]$ we add an edge from $l_{i,l}$ to c_i .

Note that the degree $d(r)$ is $2m \log n + 1$ and that all the above defined edges add to any node in a LCT at most one edge that connects it to a node outside its own LCT.

The idea of the proof is the following: The graph is built in a way that information propagation from V_ζ to V_1 is quick, but not the other way round. In the random RPULL model, wherever the rumor starts, it reaches r quickly and from there r_ζ manages to get the rumor from r in polylogarithmic time. Then the rumor quickly propagates through $V_\zeta = D_\zeta$, and from L_ζ to all LCTs D_1, \dots, D_m and afterwards to D_α .

In the adversarial RPULL model, as long as the rumor does not start in V_ζ , the rumor can quickly spread to r , a few of the D_i s and D_α but not to V_ζ because we let the adversary always prioritize a request at node r from a node in one of the D_i s over a request from r_ζ to prevent that r_ζ will get the rumor. This is possible because we show that for polynomially many rounds there is always a request at r from one of the D_i s to serve. Thus, to inform D_ζ all information must go through one of the edges $\{r_i, l_{\zeta,i}\}, i = 1, \dots, m$, with r_i informed. In less than a polynomial number of rounds few enough of the r_i s are informed and in each round only few requests from the leaf nodes L_ζ request from one of the r_i s at all making it unlikely that one of them requests from an informed r_i . Hence propagation through one of these edges is unlikely and it takes a long time for the rumor to spread over the entire graph.

Random RPULL. We start proving that random RPULL manages to spread the rumor quickly in G .

- (1) If there is an informed node in D_ζ , by **Claim 22**, w.h.p., all of D_ζ is informed in $O(\log n)$ rounds. Assume this has happened. Since each root of a LCT D_i has degree 3 in G , it requests the rumor from an informed leaf node in L_ζ w.c.p.– since no other node in L_ζ is still uninformed and therefore able to create a conflict, w.h.p., in $O(\log n)$ rounds, all root nodes r_1, \dots, r_m know the rumor.
- (2) If there is an informed node in D_i for some $i = 1, \dots, m$, due to **Claim 22**, the whole LCT D_i is informed w.h.p.(m) (=w.h.p.) within $O(\log m) = O(\log n)$ rounds. Assume this has happened. Node $c_i \in C_\alpha$ has degree at most 4 and therefore requests from its neighboring node $l_{i,m}$ w.c.p., and since all nodes in D_i are informed, it will also get the rumor.
- (3) If there is an informed node in D_α , by **Claim 22**, w.h.p., all of D_α is informed in $O(\log n)$ rounds. Assume this has happened. Almost half of all neighbors of r lie in L_α , and with same reasoning as above, r gets the rumor w.h.p. within $O(\log n)$ rounds.
- (4) Let r be informed. All its neighbors in L_α have degree $n + 1$ and therefore request with probability at most $1/n$ from r , i.e., in expectation no more than 1 node from there requests the rumor from r each turn. Each neighboring leaf node in some L_i has degree $m + 1$, i.e., requests the rumor from r with probability at most $1/m$. Since r has $m \log n$ such neighbors, in expectation no more than $\log n$ such neighbors request from r . With a Chernoff bound, w.h.p. there are no more than $O(\log n)$ requests at r . Since r_ζ has degree 3, it therefore requests w.c.p. and gets the rumor with probability $\Omega(1/\log n)$. W.h.p., the rumor is therefore propagated to r_ζ in $O(\log^2 n)$ rounds.

Altogether, wherever the source node is located, the above reasoning shows that, w.h.p., the rumor is propagated to all nodes within $O(\log^2 n)$ rounds.

Adversarial RPULL. Let $s \in V$ be the source node with the rumor. If $s \in D_i$ for some $i \in [m]$, then, without loss of generality, we assume that all nodes in D_i , D_α and r are already informed, initially. Otherwise we inform all nodes in D_α and r . For $i \in [m]$ we call any D_i *informed*, if it contains at least one informed node, otherwise *uninformed*.

The adversary has the following simple strategy. If r_ζ and at least one other node requests the rumor from r , then r chooses to pass the rumor to any other node than r_ζ . In every other aspect it follows an arbitrary strategy.

For time t we denote with X_t the number of informed LCTs D_i , and we assume without loss of generality that the corresponding LCTs are D_1, \dots, D_{X_t} . Let \mathcal{E}_t be the event that in round t no node in D_ζ has the rumor. Conditioning on this event implies that, by the structure of our graph and our model, nodes from LCTs D_i need to get the rumor from either r or from D_α , via connections $\{c_i, l_{i,m}\}$.

Let \mathcal{X}_t be the event that $X_t < 4ct + 2 \log n$, \mathcal{A}_t the event that r_ζ gets the rumor in round t and let \mathcal{C}_t be the event that a node from L_ζ gets the rumor from one of the roots r_i .

► **Claim 23.** $\mathbf{P}(\mathcal{X}_t | \mathcal{E}_t) \geq 1 - 1/n \geq e^{-1/n}$ for any $t \leq l/5c$.

Proof of Claim 23. In each round, r can inform at most one node in a yet uninformed LCT D_i . Also, any uninformed node $l_{i,l}$ connects to its neighbor c_i in D_α only with probability $1/(l+1) < 1/l$. With at most m such uninformed nodes trying to get the rumor from D_α each round, the amount of nodes $l_{i,l}$, $i = 1, \dots, m$ informed through such an edge is upper bounded by a Binomial random variable $\text{Bin}(tcl, 1/l)$. Let X'_t be the random variable that counts the number of times when an uninformed LCT D_i gets informed through such an edge to D_α but not through a connection to r . Then, by Chernoff, for $\rho = (3 - \frac{1}{c}) + 2\frac{\log n}{ct} > 1$,

$$\mathbf{P}(X_t \geq t + ct + \rho ct | \mathcal{E}_t) \leq \mathbf{P}(X'_t \geq (1 + \rho)ct) \leq \exp\left(-\frac{tc}{2}\rho \ln(1 + \rho)\right) \leq n^{-1}.$$

Therefore, w.h.p., X_t is smaller than $4ct + 2 \log n$ for $t \leq l/5c$. ◀

► **Claim 24.** $\mathbf{P}(\mathcal{A}_{t+1} | \mathcal{E}_t \cap \mathcal{X}_t) \leq 1/n \leq 1 - e^{-1/n}$ for any $t \leq l/5c$.

Proof of Claim 24. Every uninformed node $l_{i,j} \in L_i$, where $j \in [\log n]$ and $i \in [m]$, requests from r with probability $1/(l+1)$. Choosing c large enough \mathcal{X}_t implies that at least $m/2$ of the D_i s are uninformed. Thus there are at least $(m/2) \cdot \log n$ uninformed leaf nodes with a connection to r in uninformed LCTs D_i . At least $0.4c \log n$ such nodes request from r in expectation. Choosing c large enough, a simple Chernoff bound gives us that, w.h.p., at least one of these nodes requests from r . Consequently, w.h.p., r does not give the rumor to r_ζ in round $t+1$. ◀

► **Claim 25.** $\mathbf{P}(\mathcal{C}_{t+1} | \mathcal{E}_t \cap \mathcal{X}_t) \leq 1 - e^{-\frac{1}{n}(2ct + \log n)}$ for any $t \leq l/5c$.

Proof of Claim 25. By our assumption of \mathcal{E}_t , at the start of round $t+1$, no node in D_ζ has the rumor, so for \mathcal{C}_{t+1} to possibly happen, a node from L_ζ must request from one of the nodes r_1, \dots, r_{X_t} , which it does with probability $1/(n+1) < 1/n$. The probability for \mathcal{C}_{t+1} to happen is therefore

$$\mathbf{P}(\mathcal{C}_{t+1} | \mathcal{E}_t \cap \mathcal{X}_t) \leq 1 - \left(1 - \frac{1}{n}\right)^{X_t} \leq 1 - \exp\left(-\frac{1}{n}(2ct + \log n)\right).$$

We know that $\overline{\mathcal{A}_{t+1} \cup \mathcal{C}_{t+1}} \cap \mathcal{E}_t \cap \mathcal{X}_t \subseteq \mathcal{E}_{t+1}$ since if under condition \mathcal{E}_t neither \mathcal{A}_{t+1} nor \mathcal{C}_{t+1} happens, then no node in D_ζ can get informed in round $t+1$.

► **Claim 26.** $\mathbf{P}(\mathcal{E}_t) \geq e^{-\frac{1}{n}(2ct^2 + t \log n)}$ for any $t \leq l/5c$.

Proof of Claim 26. The proof follows by induction. In round $t=0$ clearly no node in D_ζ is informed, so the induction base holds. For the following, note that conditioned on \mathcal{E}_t , events \mathcal{A}_{t+1} and \mathcal{C}_{t+1} (and therefore also their complements) are independent.

$$\begin{aligned} \mathbf{P}(\mathcal{E}_{t+1}) &\geq \mathbf{P}(\overline{\mathcal{A}_{t+1} \cup \mathcal{C}_{t+1}} \cap \mathcal{E}_t \cap \mathcal{X}_t) = \mathbf{P}(\mathcal{E}_t \cap \mathcal{X}_t) \mathbf{P}(\overline{\mathcal{A}_{t+1}} \cap \overline{\mathcal{C}_{t+1}} | \mathcal{E}_t \cap \mathcal{X}_t) \\ &\geq \mathbf{P}(\mathcal{E}_t) \mathbf{P}(\mathcal{X}_t | \mathcal{E}_t) \mathbf{P}(\overline{\mathcal{A}_{t+1}} | \mathcal{E}_t \cap \mathcal{X}_t) \mathbf{P}(\overline{\mathcal{C}_{t+1}} | \mathcal{E}_t \cap \mathcal{X}_t) \\ &\geq \exp\left(-\left(\frac{1}{n}(2ct^2 + t \log n) + \frac{1}{n} + \frac{1}{n} + \frac{1}{n}(2ct + \log n)\right)\right) \\ &\geq \exp\left(-\frac{1}{n}(2c(t+1)^2 + (t+1) \log n)\right) \quad \blacktriangleleft \end{aligned}$$

This means, that after $t = \sqrt{n/c}$ rounds with probability at least e^{-3} still not all $\Theta(n)$ nodes in G are informed, concluding the proof of **Lemma 3**. ◀

C.2 Proofs for Section 3

► **Claim 17** (Quotient Rule).

$$\frac{\mu_1(A^x)}{\mu_1(A_x)} \geq \frac{\mu_2(B^x)}{\mu_2(B_x)} \text{ holds for all } A, B \in \mathcal{S}. \quad (7)$$

Proof. At first note that for any $C \in \mathcal{S}$ it holds that

$$\mathbf{P}(S_1^{\text{PULL}} = C^x | S_1^{\text{PULL}} \in (\{C^x, C_x\})) = \mathbf{P}(x \in S_1^{\text{PULL}}) = \frac{d_S(x)}{d(x)}. \quad (9)$$

This is true because in PULL x is informed independently of what else is happening and due to the fact that we already condition on S_1^{PULL} being either C^x or C_x , hence the probability of S_1^{PULL} being C^x depends solely on x being informed. Second, recall that μ_1 and μ_2 are strictly positive measures, so $\mu_1(C), \mu_2(C) > 0$ for every $C \in \mathcal{S}$, even $C = \emptyset, U$.

Let $A, B \in \mathcal{S}$ and $x \in U$. Then Lemma 14 implies

$$\begin{aligned} \Rightarrow \quad & \mathbf{P}(S^{\text{VPULL}} = A^x | S^{\text{VPULL}} \in (\{A^x, A_x\})) \geq \mathbf{P}(S_1^{\text{PULL}} = A^x | S_1^{\text{PULL}} \in (\{A^x, A_x\})) \\ & \stackrel{(9)}{=} \mathbf{P}(S_1^{\text{PULL}} = B^x | S_1^{\text{PULL}} \in (\{B^x, B_x\})) \\ \Rightarrow \quad & \frac{\mu_1(A^x)}{\mu_1(A^x) + \mu_1(A_x)} \geq \frac{\mu_2(B^x)}{\mu_2(B^x) + \mu_2(B_x)} \\ \Rightarrow \quad & \frac{\mu_1(A^x) + \mu_1(A_x)}{\mu_1(A^x)} \leq \frac{\mu_2(B^x) + \mu_2(B_x)}{\mu_2(B^x)} \\ \Rightarrow \quad & \frac{\mu_1(A_x)}{\mu_1(A^x)} \leq \frac{\mu_2(B_x)}{\mu_2(B^x)} \\ \Rightarrow \quad & \frac{\mu_1(A^x)}{\mu_1(A_x)} \geq \frac{\mu_2(B^x)}{\mu_2(B_x)}. \quad \blacktriangleleft \end{aligned}$$

► **Claim 18.** *The quotient rule (17) implies that the Holley criterion is satisfied for the measures μ_1 and μ_2 .*

Proof. ⁵

Let $A, B \in \mathcal{S}$ and $C := A \setminus B = \{c_1, \dots, c_r\}$. The Holley criterion is trivially fulfilled if $A \subseteq B$; hence assume otherwise, which implies $r \geq 1$. Write $C_s := \{c_1, \dots, c_s\}$ for $1 \leq s \leq r$. By a telescoping argument we obtain the following.

$$\frac{\mu_1(A \cup B)}{\mu_1(B)} = \frac{\mu_1(B \cup C)}{\mu_1(B \cup C_{r-1})} \cdot \frac{\mu_1(B \cup C_{r-1})}{\mu_1(B \cup C_{r-2})} \cdots \frac{\mu_1(B \cup C_1)}{\mu_1(B)} \quad (10)$$

Applying (7) to each fraction we obtain

$$\geq \frac{\mu_2((A \cap B) \cup C)}{\mu_2((A \cap B) \cup C_{r-1})} \cdot \frac{\mu_2((A \cap B) \cup C_{r-1})}{\mu_2((A \cap B) \cup C_{r-2})} \cdots \frac{\mu_2((A \cap B) \cup C_1)}{\mu_2((A \cap B))} \quad (11)$$

$$= \frac{\mu_2(A)}{\mu_2(A \cap B)}. \quad (12) \quad \blacktriangleleft$$

⁵ This proof is adapted from [14, chapter 2, page 24].