Sub-linear Universal Spatial Gossip Protocols*

Hervé Baumann[1] and Pierre Fraigniaud[1,2]

[1] University Paris Diderot
[2] CNRS
Herve.Baumann@liafa.jussieu.fr
Pierre.Fraigniaud@liafa.jussieu.fr

Abstract. *Gossip* protocols are communication protocols in which, periodically, every node of a network exchanges information with some other node chosen according to some (randomized) strategy. These protocols have recently found various types of applications for the management of distributed systems. *Spatial* gossip protocols are gossip protocols that use the underlying spatial structure of the network, in particular for achieving the "closest-first" property. This latter property states that the closer a node is to the source of a message the more likely it is to receive this message within a prescribed amount of time. Spatial gossip protocols find many applications, including the propagation of alarms in sensor networks, and the location of resources in P2P networks. We design a sub-linear spatial gossip protocol for arbitrary graphs metric. More specifically, we prove that, for any graph metric with maximum degree Δ, for any source s and any ball centered at s with size b, new information is spread from s to all nodes in the ball within $O((\sqrt{b \log b}\, \log\log b + \Delta) \log b)$ rounds, with high probability. Moreover, when applied to general metrics with uniform density, the same protocol achieves a propagation time of $O(\log^2 b \log\log b)$ rounds.

Keywords: epidemic algorithm, information spreading, resource location.

1 Introduction

Gossip protocols are communication protocols in which, periodically, every node of a network exchanges information with some other node chosen according to some (randomized) strategy. These protocols are appealing for their simplicity and robustness, and have recently found several applications for various network and system tasks, such as, e.g., multicast [5, 9], resource location [7, 8], and distributed databases management [1, 6]. In essence, a gossip protocol performs as follows. Let V be a (finite or infinite, but countable) set of nodes. The protocol is executed by all the nodes in parallel for the purpose of broadcasting information – hereafter called "gossip" – among all nodes in V. More precisely, nodes execute infinitely the same actions, in rounds, and, at each round, every node $u \in V$ applies the following two instructions:

* Both authors received additional supports from the ANR project "ALADDIN", from the INRIA project "GANG".

(1) select a node $v \in V$;
(2) send known gossips to v;

The communication between node u and the selected node v is achieved via some underlying point-to-point communication protocol which allows any pair of distinct nodes to communicate in V.

Gossip protocols differ according to (1) the way each node selects the recipient of its next point-to-point communication, and (2) the way each node chooses the gossips to be sent to that recipient. In this paper, we restrict our attention to the former point, in order to measure the impact of the node selection mechanism on the efficiency of gossip protocols. Thus, for the sake of simplicity, we assume that once the recipient v of a point-to-point communication has been selected by u, this latter node transmits to v all the gossips learnt so far. Although this assumption might be unrealistic in many contexts (because sending many gossips obviously creates congestion), it allows us to focus on the way information can spread solely as a function of the networking environment, and in absence of any hypotheses regarding the nature of the gossips. In fact, there exists several environments in which ignoring congestion created by simultaneous transmissions of many different gossips is realistic. This is typically the case of alarm spreading among nodes of a sensor network, in environments in which few nodes are expected to be simultaneously the sources of alarms (e.g., forest fires, car accidents, etc.).

Protocols that are oblivious to the past are usually preferred, for they are not sensitive to any events that occurred previously. In particular, if the protocol is oblivious, then a node recovering from a crash or a transient fault can restart the execution of the protocol from scratch, even if all local information were lost. Therefore, the selection of node v performed by an informed node u is preferably the result of a mechanism that is not depending on the past. Moreover, protocols that are also oblivious to the sources of the gossips are also preferred, for their simplicity and efficiency. In particular, by treating all sources the same, no information is required to be stored in the gossip packets, or at least the examination of this information is not required to decide to which node(s) each gossip must be forwarded.

One way to overcome the two above constraints (time obliviousness, and source independence) is to consider randomized algorithms. In fact, currently proposed practical gossip protocols [5] are based on randomized mechanisms, mostly because randomization also preserves the mechanism from possible changes in the environment.

The most popular gossip protocol is UNIFORM: at each round, every informed node u selects the recipient v uniformly at random among all nodes in V. This protocol is known to perform well in practice [5]. It has been formally analyzed by Frieze and Grimmett [4], and by Pittel [10]. In particular, the former authors have proved that a new gossip is, w.h.p., spread to all nodes in V in $O(\log n)$ rounds, where $n = |V|$. This completion time is asymptotically optimal because the number of nodes aware of a given gossip can at most double at each round, and thus it takes at least $\Omega(\log n)$ rounds for all nodes to become aware of a new gossip. However, it was noticed by Kempe, Kleinberg, and Demers [7] that UNIFORM is not appropriate to contexts in which closest nodes to the source of a new gossip should preferably receive this gossip faster than nodes farther away. Such a requirement occurs typically in the context of resource location [11] in which users are aiming at finding the nearby copies of duplicated shared resources (e.g., movies). It also occurs in the aforementioned context of alarm spreading.

Kempe et al. [7] tackled the issue of designing gossip protocols satisfying that the closer a node is from a source, the more likely it is to receive a gossip from that source within a prescribed amount of time. In order to measure the distance sensitivity of a gossip protocol, they have considered its *propagation time* as a function of the distance to the source, in a metric space (V, δ). In such a metric space, we denote by V the (finite or infinite but countable) set of points, or nodes, and by δ the distance function between nodes. Kempe et al. have designed a gossip protocol, here called DENSITY, satisfying that, if the nodes in V are spread with uniform density in the D-dimensional Euclidean space $\mathbb{R}^D$ with $\mathcal{L}_k$ metric, then a new gossip is spread to nodes at distance d from any source s in $O(\log^{1+\epsilon} d)$ rounds, with probability at least $1 - \frac{1}{\log d}$. By uniform density, it is meant that there exist two positive constants β_1 and β_2 such that, for any $r \geq 1$, the number of nodes in any ball of radius r is at least $\beta_1 r^D$, and at most $\beta_2 r^D$.

Protocols whose performances are sensitive to the distances between the sources and the recipients, are called *spatial* gossip protocols. In order to compare the propagation times of different spatial gossip protocols, in possibly different metric spaces (V, δ), we must take into account the fact that the number of nodes at a given distance from a given node varies significantly from one metric space to another, and even within the same metric space. In a metric (V, δ), it is actually more convenient to define the propagation time as a function of the ranks of the nodes, where the rank of node u relative to another node s is the number of nodes whose distances from s are not larger than $\delta(s, u)$. Indeed, a gossip protocol cannot insure that a gossip reaches a node close to the source quickly if there is a huge number of other nodes that are even closer to that source.

So let us redefine the propagation time as a function of the node ranks. The ball of radius d centered at s is defined as

$$B(s, d) = \{u \in V, \delta(s, u) \leq d\}.$$

For any node s and any $b \geq 1$, let $T_s(b)$ be the random variable equal to the number of rounds it takes for a new gossip introduced at node s to reach all nodes in the smallest ball B centered at s satisfying $|B| \geq b$. We say that a gossip protocol has propagation time $f(b)$ for some function f if for any $s \in V$, $T_s(b) \leq O(f(b))$ with high probability[3].

[3] When we write "with high probability" here, we mean with probability at least $1 - O(1/b^\alpha)$ for some $\alpha > 0$, where α may appear in the constant of the expression $O(f(b))$.

Table 1. Time complexities of various gossip protocols: n denotes the number of nodes, b the ball size, Δ the maximum degree, and D the diameter.

Protocol	Application	Propagation time	Completion time
UNIFORM	Arbitrary finite metric		$O(\log n)$ [4, 10]
DENSITY	Uniform density in $(\mathbb{R}^D, \mathcal{L}_k)$	$O(\log^{2+\epsilon} b)$ [7]	
LOCAL	Arbitrary graph metric	$O(b \log b)$	$O(\Delta(D + \log n))$ [2]
LOGSCALE	Arbitrary graph metric	$O((\sqrt{b \log b}. \log\log b + \Delta) \log b)$	
	Metric of uniform density	$O(\log^2 b \log\log b)$	

By definition of uniform density, the balls in the sub-metric $(V, \mathcal{L}_k)$ of $(\mathbb{R}^D, \mathcal{L}_k)$ induced by a set V of nodes spread out with uniform density in $\mathbb{R}^D$ have sizes polynomial in their radius. Using this fact, one can show that the gossip protocol in [7] has propagation time $O(\log^{2+\epsilon} b)$. This result yields the question of the existence of efficient gossip protocols (that is protocols with bounded propagation time) in arbitrary metrics, or at least in arbitrary graph metrics. Recall that a graph metric (V, δ) is determined by an undirected unweighted graph $G = (V, E)$, where the distance $\delta(u, v)$ between two nodes u and v is the length of a shortest path between u and v in G.

In graph metrics, a natural candidate for such a protocol is the one that uses only the links of the graph: each node selects the recipient of its next communication uniformly at random among its neighbors in the graph. We call this protocol LOCAL. This protocol has been analyzed in detail in [2], where it is shown that it completes in $O(\Delta(D + \log n))$ rounds, w.h.p., where Δ denotes the maximum degree of the nodes, and D denotes the diameter of the graph. If fact, it is not difficult to adapt results in [2] to show that the propagation time of LOCAL is $O(b \log b)$ (see Section A in the Appendix), hence proving the existence of a universal spatial gossip protocol. This bound is tight. Indeed, in the n-node star (an n-node tree with $n - 1$ leaves an one internal node called center), a gossip introduced at the center of the star will reach all nodes at distance 1 in time $\Omega(n \log n)$, by equivalence to the coupon collector problem.

The main objective of this paper is to design universal spatial gossip protocols with sub-linear propagation times.

Our results

We design a universal gossip protocol, called LOGSCALE, and prove that, in graph metrics of maximum degree Δ, its propagation time is $O((\sqrt{b \log b}. \log\log b + \Delta) \log b)$ rounds. The performances of this protocol compared to the previously mentioned protocols are summarized in Table 1. LOGSCALE has a propagation time significantly smaller than LOCAL. In finite graph metrics, it has the same completion time (i.e., the time to inform all nodes) as LOCAL, following from the fact that, in expectation, LOGSCALE acts as LOCAL for half of the rounds. During the other half, every node selects the recipient of its transmission with a probability that scales with the logarithm of the ranks. In fact, by combining

LOGSCALE with UNIFORM (every node acts as in one protocol with probability half, and as in the other protocol with probability half), we obtain a gossip protocol with the same propagation time as LOGSCALE but with the same completion time as UNIFORM. Although designed for graph metrics, our protocol LOGSCALE can also be applied to arbitrary metric. In metrics of uniform density (i.e., the same framework as in [8]), LOGSCALE achieves the polylogarithmic propagation time $O(\log^2 b \log\log b)$ rounds.

The paper is organized as follows. The gossip protocol LOGSCALE is described in Section 2, and analyzed in Section 3. The performances of LOGSCALE in metrics of uniform density are presented in Section 4. Finally, Section 5 lists some concluding remarks. (Section A in the Appendix revisits Protocol LOCAL to prove that its propagation time $O(b \log b)$ rounds.)

2 The gossip protocol LOGSCALE

This section describes the protocol LOGSCALE. The only thing one needs to specify is the way a node u selects a node v at each round. This selection process is inspired from the augmentation process in [3], in the sense that it uses a set of balls of exponentially growing size in which nodes are selected. However, as opposed to [3], the parameter k determining the size 2^k of the considered ball is not chosen uniformly at random in $[1, \log n]$, but is chosen with a probability decreasing as $1/k$. Moreover, our selection process gives high weight to neighboring nodes, as opposed to [3] which tends to ignore those neighboring nodes.

At each round, with probability $1/2$, node v is selected uniformly at random among all the neighbors of u, and, with the remaining probability $1/2$, v is selected in one ball containing 2^k nodes, for some $k > 0$. More precisely, for $k \geq 1$, let $C_k(u)$ be a set of 2^k closest nodes from u. The set $C_k(u)$ is not uniquely defined because of nodes at equal distance from u, so here $C_k(u)$ denotes one of these sets of 2^k closest nodes, chosen arbitrarily. (Note that, in finite graph metrics, $|C_k(u)| = \min\{n, 2^k\}$). To select v, node u picks one $k \geq 1$, and then selects v uniformly at random in $C_k(u)$. The choice of k is however not uniform, and k is picked with probability

$$p_k = \frac{1}{\sigma} \frac{1}{k \cdot \log^2(1+k)}.$$

where σ is a constant normalizing factor (independent of any parameter) so that $\sum_{k \geq 1} p_k = 1$. Note that choosing a larger σ would allow us to deal with transmission failures. Nevertheless, for the sake of simplicity, we assume here that $\sum_{k \geq 1} p_k = 1$. Note also that $\int_1^{+\infty} \frac{dx}{x \log^2(1+x)}$ is finite: the role of the polylog factor is specifically to insure convergence. Replacing p_k by $\frac{1}{k^{1+\epsilon}}$ would also work, but would increase the propagation time.

To sum up, let us define, for any two nodes u and v, the parameter

$$r_u(v) = \min\{k \geq 1 \mid v \in C_k(u)\},$$

and let $\Pr[u \to v]$ denotes the probability that node u selects node v at a given round. Finally, let $\deg(u)$ denote the degree of node u, i.e., the number of its adjacent nodes (neighbors). Then the protocol works as follows.

Protocol LOGSCALE**:** Set

$$p_{u,v} = \frac{1}{\sigma} \sum_{k \geq r_u(v)} \frac{1}{2^k \cdot k \cdot \log^2(1+k)}.$$

and set

$$\Pr[u \to v] = \begin{cases} \frac{1}{2}\left(\frac{1}{\deg(u)} + p_{u,v}\right) & \text{if } u \text{ and } v \text{ are neighbors,} \\ \frac{1}{2}\, p_{u,v} & \text{otherwise} \end{cases}$$

3 Propagation time of LOGSCALE

In this section, we prove our main result, namely:

Theorem 1. *For any graph metric (V, δ), and for any source node $s \in V$, protocol* LOGSCALE *satisfies that a message introduced at node s reaches all nodes in any ball of size b centered at s in less than $O((\sqrt{b \log b} \cdot \log \log b + \Delta) \log b)$ steps, with high probability.*

Proof. Let (V, δ) be a graph metric, let $s \in V$, and let B be a ball centered at s, containing $b = |B|$ nodes. We prove that, with high probability, all nodes in B receive a gossip from s, at most $O((\sqrt{b \log b}. \log \log b + \Delta) \log b)$ rounds after it appeared at s. Let $k = \lceil \log b \rceil$. We have

$$C_{k-1}(s) \subseteq B \subseteq C_k(s).$$

Let us fix $u \in C_k(s)$, and set

$$\nu = \lceil 2^{k/2} \sqrt{k} \log(1+k) \rceil.$$

For any node $x \in C_k(s)$, we define $D(x)$ as the set of the ν closest nodes from x in $C_k(s)$, where, in case of ties, node u enters $D(x)$ first. That is,

$$u \notin D(x) \Rightarrow \forall w \in D(x), \delta(s, w) < \delta(s, u).$$

Let $P = (s_0, s_1, \ldots, s_\ell)$ be a shortest path from $s_0 = s$ to $s_\ell = u$. Then let i be the smallest index such that $u \in D(s_i)$.

Claim. The expected number of rounds of LOGSCALE before s eventually selects a node $v \in D(s_i)$ is at most $2\sigma\nu$.

Proof.

$$\begin{aligned}
\Pr[s \to D(s_i)] &= \sum_{v \in D(s_i)} \Pr[s \to v] \\
&\geq \sum_{v \in D(s_i)} \frac{p_{s,v}}{2} \\
&\geq \frac{|D(s_i)|}{2\sigma} \sum_{j \geq k} \frac{1}{2^j j \log^2(1+j)} \\
&\geq \frac{\nu}{2\sigma 2^k k \log^2(1+k)} \\
&= \frac{1}{2\sigma\nu}.
\end{aligned}$$

Therefore, after an expected number of rounds $2\sigma\nu$, some node $v \in D(s_i)$ has received the gossip directly from s. This establishes the claim. $\diamond$

We now bound the expected number of rounds for the gossip to reach u from the node $v \in D(s_i)$. Let us consider the two shortest paths $P(v, s_i) = (v_0, v_1, \ldots, v_r)$ from $v_0 = v$ to $v_r = s_i$, and $P(s_i, u) = (s_i, s_{i+1}, \ldots, s_\ell)$ from s_i to $u = s_\ell$. In order to analyze the expected propagation time of the message from v to u along $P(v, s_i)$ and $P(s_i, u)$, we observe that if $d = \delta(s_i, u)$ then

$$B(s_i, d-1) \subseteq D(s_i) \subseteq B(s_i, d+1). \tag{1}$$

The first inclusion follows from the fact that $u \in D(s_i)$, hence all nodes at distance less than $d = \delta(s_i, u)$ must be in $D(s_i)$ as well by definition of the sets $D(\cdot)$. To establish the second inclusion, we first note that $D(s_{i-1}) \subseteq B(s_{i-1}, d)$ because s_{i-1} is at distance $d+1$ from u, and $u \notin D(s_{i-1})$, which implies that no other node at distance $d+1$ from s_{i-1} can be in $D(s_{i-1})$. Now, $B(s_{i-1}, d) \subseteq B(s_i, d+1)$. Hence $D(s_{i-1}) \subseteq B(s_i, d+1)$, and thus $|B(s_i, d+1) \cap C_k(s)| \geq \lceil \nu \rceil$. Therefore $D(s_i) \subseteq B(s_i, d+1)$. Let us first concentrate on the propagation time along $P(v, s_i) = (v_0, v_1, \ldots, v_r)$.

Claim. The expected number of rounds of LOGSCALE to travel from node $v \in D(s_i)$ to s_i is at most $6(\Delta + \nu)$.

Proof. We use a fact observed in [2] stating that every node outside a shortest path in a graph can be adjacent to at most 3 nodes of the path. We apply this observation in our context as follows. Any node in $D(s_i)$ can be adjacent to at most 3 nodes of $P(v, s_i)$. The problem is that some nodes outside $D(s_i)$ may also be adjacent to nodes of $P(v, s_i)$. Nevertheless, by Equation 1, only nodes $v = v_0$, v_1, and v_2 may be adjacent to nodes outside $D(s_i)$. Indeed, for all $j \geq 3$, we have $v_j \in B(s_i, d-2)$ because $v \in D(s_i) \subseteq B(s_i, d+1)$. That is v_j cannot be at the frontier between $D(s_i)$ and $V \setminus D(s_i)$ for $j \geq 3$. Therefore,

$$\sum_{j=0}^{r-1} \deg(v_j) = \deg(v_0) + \deg(v_1) + \deg(v_2) + \sum_{j=3}^{r-1} \deg(v_j)$$

$$\leq 3\Delta + 3|D(s_i)|$$
$$\leq 3(\Delta + \nu).$$

Now, the degree of a node is equal to the expected number of rounds to travel one more step along the path. As a consequence, the expected number of rounds for the gossip to travel from v to s_i is at most twice that bound (because neighbors are selected with probability $\frac{1}{2}$), and thus at most $6(\Delta + \nu)$, as claimed. $\diamond$

Let us now concentrate on the propagation time along the path $P(s_i, u) = (s_i, s_{i+1}, \ldots, s_\ell)$.

Claim. The expected number of rounds of LOGSCALE to travel from s_i to u is at most $2(3\nu + \Delta)$.

Proof. By Equation 1, all nodes s_j for $j = i, \ldots, \ell - 2$ cannot be at the frontier between $D(s_i)$ and $V \setminus D(s_i)$. As a consequence,

$$\begin{aligned} \sum_{j=i}^{\ell-1} \deg(s_j) &= \sum_{j=i}^{\ell-2} \deg(s_j) + \deg(s_{\ell-1}) \\ &\leq 3|D(s_i)| + \Delta \\ &\leq 3\nu + \Delta. \end{aligned}$$

As a consequence, the expected number of rounds for the gossip to travel from s_i to u is at most $2(3\nu + \Delta)$. $\diamond$

Let $T_{s,u}$ be random variable counting the number of round for a gossip arising at s to reach u. From what precedes, we get that

$$\begin{aligned} \mathbb{E}T_{s,u} &\leq 2\sigma\nu + 6(\Delta + \nu) + 2(3\nu + \Delta) \\ &= (2\sigma + 12)\nu + 8\Delta. \end{aligned}$$

Now, let $\alpha > 1$. For $i = 1, \ldots, \alpha \log b$, let X_i be independent random variables identically distributed as $T_{s,u}$, and denoting the time taken by a gossip starting from s at round $2(i-1)\mathbb{E}T_{s,u}$ to reach u. Since, the decision taken at each node in LOGSCALE is oblivious from the past, independent from the message source, and independent from the decision taken at other nodes, we get that if there exists i such that $X_i \leq 2\mathbb{E}T_{s,u}$, then $T_{s,u} \leq (2\alpha \log b)\, \mathbb{E}T_{s,u}$. Therefore,

$$\Pr[T_{s,u} > 2\,\alpha\, \log(b)\, \mathbb{E}T_{s,u}] \leq \prod_{i=1}^{\alpha \log b} \Pr[X_i \leq 2\mathbb{E}T_{s,u}].$$

By Markov inequality, we get $\Pr[X_i > 2\mathbb{E}X_i] < 1/2$. Therefore,

$$\Pr[T_{s,u} > 2\,\alpha\, \log(b)\, \mathbb{E}T_{s,u}] < \frac{1}{2^{\alpha \log b}} \leq \frac{1}{b^\alpha}.$$

Thus, by union-bound

$$\Pr[\exists u \in B,\ T_{s,u} > 2\,\alpha\, \log(b)\, \mathbb{E}T_{s,u}] \leq |B|\frac{1}{b^\alpha} = \frac{1}{b^{\alpha-1}}.$$

Thus

$$\Pr[\forall u \in B,\ T_{s,u} \leq 2\,\alpha\,\log(b)\,\mathbb{E}T_{s,u}] \geq 1 - \frac{1}{b^{\alpha-1}}$$

which yields

$$\Pr[\forall u \in B,\ T_{s,u} \leq 2\,\alpha\,\log(b)\,((2\sigma+12)\nu + 8\Delta)] \geq 1 - \frac{1}{b^{\alpha-1}}$$

We complete the proof by noting that

$$\nu = \left\lceil \sqrt{2^k k \log^2(1+k)} \right\rceil \leq O(\sqrt{b \log b} \log\log b).$$

□

4 Application to metrics with uniform density

Protocol LOGSCALE can also be applied to arbitrary metrics (not only graph metrics). For the protocol to run in arbitrary metrics, one simply modifies it by having the selection process defined by

$$\Pr[u \to v] = p_{u,v}$$

for any pair of nodes u, v. (There is no more condition on whether u and v are adjacent or not). We analyze LOGSCALE in the context of metrics of uniform density (cf. [7]). In this paper, we use the following definition.

Definition 1. *A metric (V, δ) has uniform density if there exists a constant c such that, for any $s \in V$, and any $k \geq 1$, we have: $\forall u, v \in V,\ u, v \in C_k(s) \Rightarrow v \in C_{k+c}(u)$.*

We shall prove that Protocol LOGSCALE performs faster in metrics with uniform density than the bound of Theorem 1. Before that, we note that Definition 1 generalizes the definition of uniform density defined in [7] for sub-merics of $(\mathbb{R}^D, \mathcal{L}_k)$. Recall that this latter definition states that a metric (V, δ) consisting of points scattered in $\mathbb{R}^D$ has uniform density if there exist two positive constants β_1 and β_2 such that, for any $r \geq 1$, the number of nodes in any ball of radius r is at least $\beta_1 r^D$, and at most $\beta_2 r^D$.

Remark. If a sub-merics of $(\mathbb{R}^D, \mathcal{L}_k)$ has uniform density in the sense of the definition in [7], then it has uniform density in the sense of Definition 1.

Proof. Assume that there exist two positive constants β_1 and β_2 such that, for any $r \geq 1$, the number of nodes in any ball of radius r is at least $\beta_1 r^D$, and at most $\beta_2 r^D$. W.l.o.g., one can assume that $\beta_1 \leq 1$. We prove that, for

$$c = 2D + \left\lfloor \log(\frac{\beta_2}{\beta_1}) \right\rfloor$$

we have: $u, v \in C_k(s)$ implies $v \in C_{k+c}(u)$ for any u, v and s.

Consider u, v and s such that $u, v \in C_k(s)$. Let r_{min} be the smallest radius such that $C_k(s) \subseteq B(s, r_{min})$.

We first analyze the "general" case $r_{min} > 1$. Let $1 \leq r' < r_{min}$ be such that $r_{min} \leq r' + 1$. We have $|B(s, r')| < 2^k$ because $r' < r_{min}$. On the other hand, we have $\beta_1 r'^D \leq |B(s, r')|$ by uniform density (in the sense of [7]). Thus $r' \leq \left(\frac{2^k}{\beta_1}\right)^{1/D}$. Now, by the triangle inequality, $\delta(u, v) \leq 2r_{min}$. Thus we get:

$$\begin{aligned}
|B(s, \delta(u, v))| &\leq |B(s, 2r_{min})| \\
&\leq \beta_2 \, (2r_{min})^D \\
&\leq \beta_2 \, 2^D (r' + 1)^D \\
&\leq \beta_2 \, 2^D \left(\left(\frac{2^k}{\beta_1} \right)^{1/D} + 1 \right)^D \\
&\leq \beta_2 \, 2^{2D} \frac{2^k}{\beta_1}.
\end{aligned}$$

Combining this latter inequality with the fact that $v \in B(s, \delta(u, v))$, we get that

$$v \in C_{\lfloor \log(\beta_2 \; 2^{2D} \frac{2^k}{\beta_1}) \rfloor}(u).$$

That is, $v \in C_{k+c}(u)$.

The "particular" case $r_{min} \leq 1$ can be treated similarly. We have

$$|B(s, \delta(u, v))| \leq |B(s, 2r_{min})| \leq \beta_2 \, 2^D.$$

Thus $v \in C_{D+\lfloor \log(\beta_2) \rfloor}(u) \subseteq C_c(u) \subseteq C_{c+k}(u)$. □

Theorem 2. *For any metric (V, δ) with uniform density, and for any source node $s \in V$, protocol* LOGSCALE *satisfies that a message introduced at node s reaches all nodes in any ball of size b centered at s in less than $O(\log^2 b \log \log b)$ steps, with high probability.*

Proof. The proof follows the same guidelines as the analysis of protocol DENSITY in [7]. Let $s \in V$, and $t \in C_k(s)$ for some k such that $k - 2 \log k \geq 2c$ where c is the constant appearing in Definition 1. Let

$$U_s = C_{\lceil k/2 + \log k \rceil + c}(s) \text{ and } U_t = C_{\lfloor k/2 + \log k \rfloor + c}(t).$$

Fix $\beta > 2$, and let $\mathcal{E}$ denotes the event "there exists at least one call from U_s to U_t occurring during at least one of βk consecutive steps".

Claim. We have $\Pr[\mathcal{E}] \geq 1 - \frac{1}{2^{\beta k}}$.

Proof. By the uniform density hypothesis, we have $C_k(s) \subseteq C_{k+c}(t)$. Therefore, $U_s \subseteq C_{k+c}(t)$. Therefore, again by the density hypothesis, we get that, for any $u \in U_s$, and any $v \in C_{k+c}(t)$, $v \in C_{k+2c}(u)$. Thus, in particular,

$$\forall u \in U_s, \; \forall v \in U_t, \; v \in C_{k+2c}(u).$$

Using that property, one can bound the probability p of a call from U_s to U_t. We have

$$1-p = \prod_{u\in U_s, v\in U_t} (1-p_{u,v}).$$

As in the proof of Theorem 1, one can easily check that $p_{u,v} \geq x$ where

$$x = \frac{1}{2\sigma 2^{k+2c}(k+2c)\log^2(1+k+2c)}.$$

Therefore

$$1-p \leq (1-x)^{|U_s|\,|U_t|} \leq e^{-x\,|U_s|\,|U_t|}$$

and thus

$$1-p \leq e^{-k^2/(2\sigma(k+2c)\log^2(1+k+2c))}.$$

For k big enough, we get that $1-p \leq 1/2$. As a consequence,

$$\Pr[\mathcal{E}] \geq 1 - \frac{1}{2^{\beta k}}$$

as claimed. $\diamond$

Now, we prove the following claim:

Claim. For any $t \in C_k(s)$, t receives a message originated at s in time at most $\beta\, k\, g(k)$ with probability at least $1-\frac{g(k)}{2^{\beta k}}$ where $g(k)$ is solution of the recurrence equation

$$g(k) = 1 + g(\lceil k/2 + \log k \rceil + c) + g(\lfloor k/2 + \log k \rfloor + 2c).$$

Proof. We establish the claim by induction. We consider three consecutive time intervals:

$$\begin{aligned} I_s &= [1\ ,\ \beta\, k\, g(\lceil k/2 + \log k\rceil + c)] \\ I &= [|I_s| + 1\ ,\ |I_s| + \beta\, k] \\ I_t &= [|I_s| + |I| + 1\ ,\ |I_s| + |I| + \beta\, k\, g(\lfloor k/2 + \log k \rfloor + 2c)]. \end{aligned}$$

By the bound we have previously derived on the event $\mathcal{E}$, we get that during the time interval I, there exists a node $u \in U_s$ which calls a node in $v \in U_t$ with probability at least $1-\frac{1}{2^{\beta k}}$. By induction hypothesis, node u has received the message of source s during time interval I_s with probability $1-\frac{g(\lceil k/2+\log k\rceil + c)}{2^{\beta k}}$. Also, by induction hypothesis, as $t \in C_{\lfloor k/2+\log k\rfloor+2c}(v)$, node t has received the message of source v during time interval I_t with probability $1-\frac{g(\lfloor k/2+\log k\rfloor+2c)}{2^{\beta k}}$. Therefore during a time interval of duration $|I_s|+|I|+|I_t| = \beta\, k\, g(k)$ node t has received a message of source s with probability

$$(1-\frac{1}{2^{\beta k}})(1-\frac{g(\lceil k/2+\log k\rceil + c)}{2^{\beta k}})(1-\frac{g(\lfloor k/2+\log k\rfloor + 2c)}{2^{\beta k}})$$

and thus with probability at least $1-\frac{g(k)}{2^{\beta k}}$. $\diamond$

Now, it is easy to see that $g(k) \leq O(k \log k)$. Thus, applying the claim, we get that for any $t \in C_k(s)$, t receives a message originated at s in time at most $O(\beta k^2 \log k)$ with probability at least $1 - O(\frac{k \log k}{2^{\beta k}})$. The theorem follows by applying union bound on all t's in $C_k(s)$. $\square$

5 Conclusion

In this paper, we have proved that there exists a universal gossip protocol for all graph metric, whose propagation time is $O((\sqrt{b \log b} \cdot \log\log b + \Delta) \log b)$. A natural question is whether this bound can be improved. In particular it would be quite informative to prove or disprove the existence of a universal gossip protocol with polylogarithmic propagation time $O(\log^\alpha b)$ for some $\alpha \geq 1$. Such a polylogarithmic propagation time can be achieved in specific metrics, namely those with uniform density. Another natural extension of this work would thus be to extend the result to arbitrary metric without any assumption on the density. In fact, even the existence of a gossip protocol with finite propagation time is not clear in this general context: UNIFORM has an unbounded propagation time, and DENSITY and LOGSCALE have polylogarithmic propagation times in metrics with uniform density, but their performances in arbitrary metrics are not known.

6 Acknowledgements

All the authors are thankful to George Giakkoupis for helpful discussions.

References

1. A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, D. Terry. Epidemic Algorithms for Replicated Database Maintenance. In 6th ACM Symposium on Principles of Distributed Computing (PODC), pages 1-12, 1987.
2. U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. In Int. Symposium on Algorithms (SIGAL), LNCS 450, pp 128-137, Springer, 1990.
3. P. Fraigniaud, C. Gavoille, A. Kosowski, E. Lebhar, and Z. Lotker. Universal Augmentation Schemes for Network Navigability: Overcoming the $\sqrt{n}$-Barrier. In 19th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), pages 1-7, 2007.
4. A. Frieze and G. Grimmett. The shortest-path problem for graphs with random arc-lengths. Discrete Applied Math. 10:57-77, 1985.
5. I. Gupta, A.-M. Kermarrec, A. Ganesh. Efficient Epidemic-Style Protocols for Reliable and Scalable Multicast. In 21st Symposium on Reliable Distributed Systems (SRDS), pages 180-189, 2002.
6. D. Kempe, A. Dobra, and J. Gehrke. Computing Aggregate Information using Gossip. In 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2003.
7. D. Kempe, J. Kleinberg, and A. Demers. Spatial gossip and resource location protocols, In 33rd ACM Symposium on Theory of Computing, pages 163-172, 2001.
8. D. Kempe and J. Kleinberg. Protocols and impossibility results for gossip-based communication mechanisms. In Proc. 43st IEEE Symp. on Foundations of Computer Science, pages 471-480, 2002.
9. J. Luo, P. Eugster, J.-P. Hubaux. Route driven gossip: probabilistic reliable multicast in ad hoc networks. In 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), pages 2229- 2239, 2003.

10. B. Pittel. On spreading a rumour. SIAM J. Applied Math., 47:213-223, 1987.
11. G. Plaxton, R. Rajaraman, and A. Richa. Accessing Nearby Copies of Replicated Objects in a Distributed Environment. In 9th ACM Symposium on Parallel Algorithms and Architectures (SPAA), pages 311-320, 1997.

APPENDIX

A Protocol LOCAL revisited

In this section, we briefly revisit the protocol LOCAL analyzed in [2], and prove that its propagation time is $O(b \log b)$.

Proposition 1. *For any finite graph metric (V, δ), and for any source node $s \in V$, protocol* LOCAL *satisfies that a message introduced at node s reaches all nodes in any ball of size b centered at s in less than $O(b \log b)$ steps, with high probability.*

Proof. We use the same proof structure as in [2]. Let (V, δ) be a finite graph metric, let $s \in V$, and let B be a ball centered at s, containing $b = |B|$ nodes. We prove that, using LOCAL, all nodes in B receive a gossip from s at most $O(b \log b)$ rounds after it appeared at s, with high probability. For this, we use again the observation in [2] stating that every node outside a shortest path in a graph can be adjacent to at most 3 nodes of the path. Let $u \in B$, and $P = (u_0, u_1, \ldots, u_\ell)$ be a shortest path from s to u, with $u_0 = s$ and $u_\ell = u$. Any node in $B \setminus P$ can be adjacent to at most 3 nodes of P. A node outside B adjacent to P can only be adjacent to u_ℓ since otherwise it would be in B. Therefore

$$\sum_{i=0}^{\ell-1} \deg(u_i) \leq 3b.$$

From this bound, we get that if X_u denotes the random variable equal to the time it takes for a gossip to reach u, then $\mathbb{E}X_u \leq 3b$. Now, let $\alpha > 1$. For $i = 1, \ldots, \alpha \log b$, let Y_i be independent random variables identically distributed as X_u. Since, the decision taken at each node in LOCAL is oblivious from the past, independent from the message source, and independent from the decision taken at other nodes, we have

$$\Pr[X_u \geq 2\ \alpha\ \log(b)\ \mathbb{E}X_u] \leq \Pi_{i=1}^{\log b} \Pr[Y_i \geq 2\mathbb{E}X_u].$$

By Markov inequality, we get

$$\Pr[X_u \geq 2\ \alpha\ \log(b)\ \mathbb{E}X_u] \leq \left(\frac{1}{2}\right)^{\alpha \log b} = \frac{1}{b^\alpha}.$$

Thus $\Pr[X_u \geq 6\ \alpha\ b\ \log b] \leq 1/b^\alpha$. By union-bound, we get that

$$\Pr[\exists u \in B, X_u \geq 6\ \alpha\ b\ \log b] \leq \frac{1}{b^{\alpha-1}}$$

which completes the proof. □