



Infective flooding in low-duty-cycle networks, properties and bounds[☆]

Luca Baldesi^{a,*}, Leonardo Maccari^b, Renato Lo Cigno^c

^a Department of Information Engineering and Computer Science, University of Trento, Italy

^b Department of Environment and Computer Science and Statistics, University of Venice Ca' Foscari, Italy

^c Department of Information Engineering, University of Brescia, Italy

ARTICLE INFO

Keywords:

Low-duty-cycle networks
Flooding
Delay
Mesh networks
Unstructured
Infective models
Optimal dissemination

ABSTRACT

Flooding information is an important function in many networking applications. In some networks, as wireless sensor networks or some ad-hoc networks it is so essential as to dominate the performance of the entire system. Exploiting some recent results based on the distributed computation of the eigenvector centrality of nodes in the network graph and classical dynamic diffusion models on graphs, this paper derives a novel theoretical framework for efficient resource allocation to flood information in mesh networks with low duty-cycling without the need to build a distribution tree or any other distribution overlay. Furthermore, the method requires only local computations based on each node neighborhood. The model provides lower and upper stochastic bounds on the flooding delay averages on all possible sources with high probability. We show that the lower bound is very close to the theoretical optimum. A simulation-based implementation allows the study of specific topologies and graph models as well as scheduling heuristics and packet losses. Simulation experiments show that simple protocols based on our resource allocation strategy can easily achieve results that are very close to the theoretical minimum obtained building optimized overlays on the network.

1. Introduction

Flooding, the function of sending a piece of information to all nodes in a network, is a fundamental and pervasive function in many protocols, applications, and network architectures as well. Flooding of Link-State (LS) advertisements in LS routing protocols or streaming in a multicast group with Peer-to-Peer (P2P) technologies are examples of flooding in application overlays, [2–5]. In wireless ad-hoc networks such as WSNs it is normally executed on the physical topology (as opposed to a logical overlay) and it is so important that its performance impacts the overall network efficiency [6–8]. In these networks, which are at the base of Internet of Things (IoT) [9,10], flooding pertains to sensor data, queries, or messages about diagnosis, localization, routing, and configuration: In practice in every domain of operation [11].

Flooding can be often solved satisfactorily in traditional networks and overlays with techniques that build a distribution tree [12,13] or similar structures, or with brute-force approaches such as limited flooding (as in Open Shortest-Path First (OSPF) LS advertisements), in WSN there are three additional challenges: (i) Dynamism; (ii) Energy consumption; and (iii) Duty cycling, i.e., the ratio between the wake-up time of the node and the overall time of the cycle [14] that can be as low as 0.01 or even less [7] when high energy efficiency is

required. A WSN is dynamic, meaning that even if the nodes are stationary, the surrounding conditions vary and every now and then we expect a few links to appear or disappear. Energy efficiency is hampered by continuous signaling, reduction in duty cycling, need for overhearing messages. Low duty cycle means that any reconfiguration takes a long time as nodes seldom wake up, but also that broadcast at the physical layer cannot be exploited, because the duty-cycling is such that only one pair of node at a time can communicate. This may look weird, but considering that overhearing messages costs a lot of energy, ensuring that only the intended destination hears packets is often the best solution [15].

For these reason we consider a flooding strategy that does not rely on trees, that are intrinsically fragile even in presence of minimal modifications of the topology, and that is intrinsically based on “cycles” by design, which maps perfectly with low duty-cycling networks.

In the rest of the paper we concentrate, as reference scenario, on WSNs, albeit our results are general and apply to any network, physical or logical, provided there are reasons to avoid wasting resources and a notion of communication cycle is present. In this context we use the term *flooding* to identify the distribution of information to all nodes in a network, while the term *broadcast* is used only to refer to physical

[☆] This research is supported by the European Commission, H2020-ICT-2015 Programme, Grant Number 688768 ‘netCommons’ (Network Infrastructure as Commons). A preliminary version of this paper has been presented at IFIP/IEEE WONS 2019 [1]. This work was done when all the authors were with the University of Trento.

* Corresponding author.

E-mail addresses: luca@baldesi.ovh (L. Baldesi), leonardo.maccari@unive.it (L. Maccari), renato.locigno@unibs.it (R. Lo Cigno).

layer broadcast, even if we never use physical broadcast properties in this paper.

The contribution of this paper stems from the encounter of a recent result published in the context of P2P streaming [16], with classical results from epidemic diffusion based on differential equations [17,18]. We look at the flooding process as the “infective” propagation of data on a graph: Each node that has already received a message can infect his own neighbors by sending that message. This observation can indeed be applied to any flooding technique or protocol. Exploiting the results in [16] we propose to modify the infective capability, i.e., the amount of information sent per cycle, of each node in such a way to optimize the information distribution. We call this modification RE, since it is based on the imposition that each node in the network, regardless of its position in the topology, has the same average probability of receiving the information. The result is a sound theoretical framework providing upper and lower stochastic bounds for the maximum flooding delay. Flooding is performed without building a distribution tree or any other structure that requires signaling or global coordination, thus resulting in a very robust system that requires minimal signaling and adapts naturally to topology changes. Furthermore, we constrain the total consumption of transmission resources by all nodes to remain constant, so that the gains provided are due only to better use of resources and not to the use of more resources.

The theoretical results we derive are confirmed and validated by event driven simulations on different topologies, applying scheduling heuristics to improve performance, and in presence of packet losses. This paper extends and completes the work presented at WONS 2019 [1] providing validation and performance results on RE-based flooding in different topology types, for networks up to 2000 nodes, and checking that packet scheduling heuristics as well as packet losses do not hamper the theoretical properties of RE-based flooding and do not invalidate the stochastic bounds, which are obtained abstracting from any specific topology or technology. The insights gained with this work and the implementation simplicity of RE allocation strategies, which do not require any centralized computation, or heavy signaling, open up the possibility of designing extremely efficient flooding protocols.

The rest of the paper is organized as follows: we present our infection flooding model in Section 2; the main theoretical results, the stochastic delay bounds, are derived in Section 3, in Section 4 we present simulation results validating our framework, Section 5 describes related works and, finally, future works and conclusions are detailed in Section 6.

2. System model

We consider a connected, multi-hop network described by an undirected graph $G(V, E)$, where V is the set of nodes and $E \subset (V \times V)$ is the set of edges. The network is stationary or with slow mobility as assumed also in [7,15,19], so that in general the network topology does not change too much from one wake-up cycle to the next. One node, called source, is the originator of the message to be flooded, but we do not make any further assumption on its location in G ; it is likely that the source changes from one flooding event to another. This scenario is typical of a WSN used for monitoring, in which at a certain instant a sensor detects a certain event and alerts the other nodes of the network. Similarly, flooding is needed for time synchronization of nodes, which can be triggered by any of the sensors acquiring time from an external source [20]. In general, any networking application that needs to obtain a distributed consensus requires, at least from time to time, to perform flooding of information.

As we observed in the introduction, tree-based overlays are commonly built on the mesh network to perform flooding operations; however, in tree-based flooding, if the source changes the flooding tree must be recomputed, e.g., using Dijkstra algorithm, unless a global Minimum Spanning Tree is used, which however does not minimize the distribution delay, and is very sensitive to topology changes. Moreover,

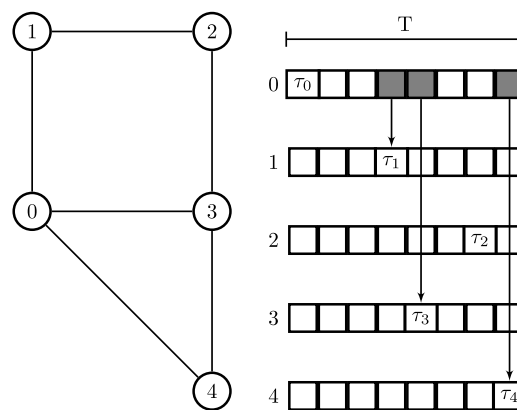


Fig. 1. The network model consists of a mesh network whose nodes have periodic listening periods τ_i and wake up when necessary in period τ_j to transmit to node j ; transmissions are unicast as τ_i are separated; in the picture node 0 sends a packet (in gray) to nodes 1, 3, and 4.

the time interval between two flooding events can be orders of magnitude larger than the wake-up cycle, so changes in the topology are likely between two flooding events. Tree structures are fragile and must be maintained over time, requiring a very large signaling overhead for low duty-cycle networks.

Let T be the wake-up cycle. T is divided in fixed lengths time slots of τ seconds; the time slot must be long enough to allow the sender and the receiver to synchronize transmit a packet even in presence of clock drifts. The ratio between T and τ also relates to the number of nodes that can be in the network without generating pathological conflicts in the access. For the sake of analytic tractability we assume that transmissions are successful, as our focus is on the assignment of transmission resources to nodes based on their topological position to minimize the flooding delay, so re-transmissions are not essential to the problem. This assumption is relaxed in Section 4.4, where, via simulations, we show that losses do not change the ranking of flooding strategies, and the one we propose remains the best performing one. Indeed, the presence of losses do not affect the flooding strategy and resource allocation proposed in this paper, neither the infective flooding model; however, they make the bounds derived in Section 3 analytically untractable.

Each node wakes up for a single τ slot during a cycle T to listen for incoming packets and sleeps during the others to reduce energy consumption, interference, avoid overhearing, or for any other reason. This assumption is the same adopted, for instance, in [19]. We assume there is an initialization phase (how it is done is outside the scope of this paper, see for instance the solution proposed in [21]) during which nodes select their listening periods and exchange them with their neighbors. After the initialization, nodes wake up if they have to send packets during other nodes listening periods or if they have a scheduled listening period.

Fig. 1 depicts the situation with 5 nodes and a simple topology. τ_i identifies the listening slot of duration τ selected by node i . As the duty-cycle is very low, one can assume that waking periods are different in the entire network, minimizing interference, but it is enough that they are different in 2-hop neighborhoods. In Fig. 1, node 0 has some information to flood. Node 0 has three neighbors, namely 1, 3 and 4, so that, during a time cycle T , it can send the information to any possible subset of its neighbors; if it sends it to all of them it reduces the reception delay, however, it consumes more resources that can possibly be wasted in the case its neighbors already received the information from other nodes (this can happen in larger networks with multi-paths).

We consider each time cycle T as a unit of time during which each node owning information can take the decision to send it to one or more of its neighbor (a function called *scheduling*).

In the following subsections we combine results on the optimality of receiver-equal resource allocation [16] together with an infective model dissemination on graphs to derive bounds on the performance of information flooding on mesh networks with minimal signaling.

Unlike the widely used flooding strategy for which each node uses the same amount of resources per unit of time [22], which we call SE, a RE strategy guarantees that in a streaming application, in which a source injects into the network a sequence of packets, every node receives the same amount of information at steady state given that the total use of resources are minimal. Minimal resources means that, if B_s bit/s is the stream bandwidth, then the overall capacity allocated in the network is $|V| \times B_s$. This node resource allocation is proportional to the eigenvector centrality of the nodes in $G(V, E)$.

The eigenvector centrality of a network node measures its *importance* with respect to its neighbor importance. Such interpretation has been used for evaluating node inter-influence and overall impact [23] for example in the Google PageRank application [24].

The RE strategy in streaming guarantees the sustainability of streaming for every node, because every node receive the same amount of information, equal to the stream rate, per time unit. Other strategies either use more resources or imply that some node receives less than the stream rate, thus lagging behind the others, steadily increasing the playout delay. This same property implies that a RE strategy in flooding minimizes the overall flooding delay, i.e., the time when all the nodes have received the information.

An infective model is, in some sense, the mathematical formulation of opportunistic (or stochastic) flooding. Each node that possesses the information pass it to one of its neighbors following some protocol, until all his neighbors have received the information. When this is true for all nodes, then the information has been flooded to the entire network.

The goal of this paper is to show that the steady-state reception-equal property can be applied to flooding of one packet and derive time bounds for the average case that are independent of the mesh network topology, so that they can be used as general indications to dimension applications and/or network resources.

2.1. Reception-equal P2P streaming

In a P2P streaming application there is a source of the stream that injects packets in the network. Every node (including the source) at every instant has a buffer of packets that can share with its neighbors, and decides which packet to share with which neighbor. As stated above, at steady state, the optimal resource allocation is the one that guarantees that every node receives the same amount of information (the same number of packets) per time-interval. The result presented in [16] can be summarized as follows.

Let A' be a stochastic transition matrix for a network $G(V, E)$ as described in Section 2, so that the element $A'_{ij} \in [0, 1]$, $A'_{ij} > 0 \iff (i, j) \in E$ represents the probability for node j to send a packet to node i during a cycle T and $\bar{1}^T A' = \bar{1}^T$. A' can be regarded as the adjacency matrix of G whose values represent transmission probabilities.

Let θ_j be the throughput (in terms of packets sent per time cycle T) that node j sustains on average and θ the resulting column vector. Then, from Theorem 1 in [16] and assuming minimal resource usage we have:

$$\theta_j = x_j \sum_{k=1}^{|V|} \frac{A'_{kj}}{x_k}, \quad A_{ij} = \frac{A'_{ij}}{x_i} \quad (1)$$

$$\bar{1} = A\theta \quad (2)$$

$$|\theta| = |\bar{1}| = |V| \quad (3)$$

$$\bar{1}^T A = \bar{1}^T \quad (4)$$

$$A'_{ij} = 0 \iff A_{ij} = 0 \quad (5)$$

where $x_i \in \mathbb{R}$ is the eigenvector centrality of node i . Theorem 1 in [16] states that the new stochastic transition matrix A describes the same links as A' but with different values (Eqs. (4) and (5)), which describe the probability to send a packet to the neighbors. θ_j represents the number of packets node j sends during T and averaging over all nodes exactly one packet per time cycle is sent (Eq. (3)), which guarantees that overall the transmission resources remain the same. Eq. (2) ensures that every node has the same probability of receiving the information if we average over all possible sources $s \in V$.

In distributed systems A' represents local strategies for forwarding. A general and wide-spread heuristics is to give the same transmission probability to every neighbor [22]; this in turns makes A' to be column-uniform (i.e., each column contains either 0 or a column-specific constant). If A' is column-uniform, which means that packets are sent with uniform probability to the neighbors, which is a very reasonable assumption, these parameters can be computed locally by each node simply gossiping their neighborhood set size [16], a property that guarantees a very simple and low overhead implementation even in very large networks. In general, the computation of the eigenvector centrality may be complex, however, as proven in [25], it can be computed with a distributed algorithm. Section 3 derives stochastic upper and lower delay bounds independent of the network topology starting from this elegant result.

In presenting results, for the sake of comparison, we consider also the SE strategy where each node sends the same amount of information at every time cycle. To the best of our knowledge, however, it is not possible to obtain bounds as those derived in Section 2.2 for the SE strategy; although only a conjecture, this may indicate that for SE strategies the actual delay bounds for flooding without any topological constraint and minimal resource use do not exist.

2.2. The infective model

The flooding of a packet in a network can be seen as a virus propagation starting at the source node, and all nodes being susceptible to the infection while they do not have the packet and infective when they have it. We are interested in studying and characterizing the speed of such infection.

Our infection process corresponds to the elementary SI model: a node can be in either one of the two states, susceptible (S) or infected (I), there is no recovery from the infection and nodes remain infectious indefinitely (they do not die or recover from the infection) [17]. In networking terminology this means that nodes that have the information continue to distribute it until all its neighbors have it, ensuring flooding. We are aware that there is a large body of literature on disease spreading, obviously in the medical literature, but also in networking (see for instance [18,26–29] and references in these works), but indeed this simple SI model represents exactly what happens flooding a packet into a network, taking into account the topological properties of the network graph G .

The initial spread of a virus in a network subject to the SI model is exponentially fast [17] and it depends on the largest eigenvalue of A' (1 in our case as it is column stochastic) and the rate of infection. During this initial phase, the nodes with large eigenvector centralities are more likely to be infected [17].

We call $y_i(k)$ the probability that node i is infected at time k (we use discrete time to better map the time cycle T). $S(k)$, $I(k)$ are the group of susceptible and infected nodes at time k , and N_i is the set of neighbor nodes of i . Hence, the following dynamic equation holds:

$$y_i(k+1) = y_i(k) + P\{i \in S(k), j \in I(k), j \text{ infects } i, \forall j \in N_i\} \quad (6)$$

Eq. (6) states that the probability that node i is infected at time $k+1$ is given by the same probability at the previous time step plus the probability of transition from the susceptible state to the infected one, which occurs if at least one neighbor j is infected (at time k) and pass the infection.

Unfortunately, Eq. (6) cannot be integrated in closed form conditioned on the graph topology, and it is hence difficult to handle mathematically. To ease the analysis we take advantage of its first order approximation (see Section 3.3 for a discussion on this approximation):

$$y_i(k+1) = y_i(k) + (1 - y_i(k))(1 - P\{j \in I(k), \forall j \in N_i, j \text{ does not infect } i\}) \quad (7)$$

In the case of packets flooding using the reception-equal strategy we have that j infects i (j sends a packet to i) with probability $A_{ij}\Theta_j$ (the throughput of j multiplied by the probability of sending a packet to neighbor i). Eq. (7) can be expressed in closed form as:

$$y_i(k+1) = y_i(k) + [1 - y_i(k)] \left[1 - \prod_{j=1}^{|V|} [1 - y_j(k)A_{ij}\Theta_j] \right] \quad (8)$$

the next section exploits this approximation to derive closed form upper and lower stochastic delay bounds.

3. Flooding delay bounds

To derive the upper and lower stochastic bounds for the distribution delay, we assume that each node in a network can be the source with uniform probability, i.e., $y_i(0) = \frac{1}{|V|} \forall i$.

We first state the bounds formulation, and then we prove, in **Theorem 1**, that they limit, in the stochastic sense, the evolution of the SI diffusion model $y_i(k)$. The upper bound is:

$$\begin{cases} \omega(k+1) &= 2\omega(k) - \omega^2(k) \\ \omega(0) &= \frac{1}{|V|} \end{cases} \quad (9)$$

and the lower bound is:

$$\begin{cases} \Omega(k+1) &= 2\Omega(k) - \frac{3}{2}\Omega^2(k) + \frac{\Omega^3(k)}{2} \\ \Omega(0) &= \frac{1}{|V|} \end{cases} \quad (10)$$

Both Eqs. (9) and (10) have two fixed points $\{0, 1\}$ the latter of which is attractive; $\omega(k), \Omega(k)$ are monotonically increasing functions and, given their initial value $\omega(0) = \Omega(0) = \frac{1}{|V|}$ their values are in the interval $[\frac{1}{|V|}, 1)$. Moreover, $\omega(k) \geq \Omega(k), \forall k$.

Theorem 1. Given a uniform initial probability $y_i(0) = \frac{1}{|V|}, \forall i$, then

$$\Omega(k) \leq y_i(k) \leq \omega(k), \forall i, k$$

Proof. Given the reception-equal property (Eq. (2)) the following identities hold for any k :

$$\begin{aligned} \omega(k+1) &= 2\omega(k) - \omega^2(k) = \\ &\omega(k) + (1 - \omega(k)) \sum_{j=1}^{|V|} \omega(k)A_{ij}\Theta_j = \\ &\omega(k) + (1 - \omega(k)) \left[1 - \left(1 - \sum_{j=1}^{|V|} \omega(k)A_{ij}\Theta_j \right) \right] \end{aligned} \quad (11)$$

and

$$\begin{aligned} \Omega(k+1) &= 2\Omega(k) - \frac{3}{2}\Omega^2(k) + \frac{\Omega^3(k)}{2} = \\ &\Omega(k) + (1 - \Omega(k)) \left[\sum_{j=1}^{|V|} \Omega(k)A_{ij}\Theta_j + \right. \\ &\left. - \frac{1}{2} \sum_{j=1}^{|V|} \sum_{z=1}^{|V|} \Omega(k)A_{ij}\Theta_j \Omega(k)A_{iz}\Theta_z \right] = \\ &\Omega(k) + [1 - \Omega(k)] \left[1 - \left(1 - \sum_{j=1}^{|V|} \Omega(k)A_{ij}\Theta_j + \right. \right. \end{aligned}$$

$$\left. \left. - \frac{1}{2} \sum_{j=1}^{|V|} \sum_{z=1}^{|V|} \Omega(k)A_{ij}\Theta_j \Omega(k)A_{iz}\Theta_z \right) \right] \quad (12)$$

The last element of Eq. (12) is represented compactly as

$$\frac{1}{2} \sum_{j=1}^{|V|} \sum_{z=1}^{|V|} b_j b_z$$

where

$$b_j = \Omega(k)A_{ij}\Theta_j \in \mathbb{R}^+.$$

Since

$$\frac{1}{2} \sum_{j=1}^{|V|} \sum_{z=1}^{|V|} b_j b_z = \sum_{j=1}^{|V|} \sum_{z>j}^{|V|} b_j b_z + \frac{1}{2} \sum_{j=1}^{|V|} b_j^2$$

then

$$\frac{1}{2} \sum_{j=1}^{|V|} \sum_{z=1}^{|V|} b_j b_z \geq \sum_{j=1}^{|V|} \sum_{z>j}^{|V|} b_j b_z$$

and finally

$$\begin{aligned} \Omega(k+1) &\leq \Omega(k) + \\ &[1 - \Omega(k)] \left[1 - \left(1 - \sum_{j=1}^{|V|} \Omega(k)A_{ij}\Theta_j + \right. \right. \\ &\left. \left. \sum_{j=1}^{|V|} \sum_{z>j}^{|V|} \Omega(k)A_{ij}\Theta_j \Omega(k)A_{iz}\Theta_z \right) \right] \end{aligned} \quad (13)$$

The derivation of the bounds is now by induction over k .

Upper bound.

We first prove $y_i(k) \leq \omega(k)$ by induction over k .

($k = 0$) $\omega(0) \geq y_i(0) \forall i$ by definition.

Assuming for the induction $\omega(k) \geq y_i(k) \forall i, k = 1, \dots, z$

($k = z + 1$) From Eq. (2) we have that $A_{ij}\Theta_j \leq 1 \forall i, j$, and the following three inequalities hold:

$$y_i(k)A_{ij}\Theta_j < 1; \Omega(k)A_{ij}\Theta_j < 1; \omega(k)A_{ij}\Theta_j < 1 \quad (14)$$

and, thus, we can apply the Weierstrass product inequality,

$$\prod_{i=1}^n (1 - a_i) \geq 1 - \sum_{i=1}^n a_i$$

(with $a_j = \omega(z)A_{ij}\Theta_j$) to Eq. (11), obtaining the following:

$$\omega(z+1) \geq \omega(z) + (1 - \omega(z)) \left[1 - \prod_{j=1}^{|V|} (1 - \omega(z)A_{ij}\Theta_j) \right]$$

For simplicity we call $\psi = 1 - \prod_{j=1}^{|V|} (1 - \omega(z)A_{ij}\Theta_j)$ then,

$$\begin{aligned} y_i(z+1) &= y_i(z) + (1 - y_i(z)) \left[1 - \prod_{j=1}^{|V|} (1 - y_j(z)A_{ij}\Theta_j) \right] \\ &\leq y_i(z) + (1 - y_i(z))\psi \end{aligned}$$

as, because of the inductive step, $\omega(z) \geq y_i(z) \forall i, z$.

Subtracting $y_i(z+1)$ from $\omega(z+1)$ we get

$$\begin{aligned} \omega(z+1) - y_i(z+1) &\geq \omega(z) - y_i(z) + (y_i(z) - \omega(z))\psi \\ &= (\omega(z) - y_i(z))(1 - \psi) \end{aligned}$$

as $\omega(z) \geq y_i(z)$ and $\psi < 1$, then $(\omega(z) - y_i(z))(1 - \psi) \geq 0$ and finally

$$\omega(z+1) \geq y_i(z+1).$$

Lower bound.

The proof of $y_i(k) \geq \Omega(k)$ is again by induction over k ,

($k = 0$) $\Omega(0) \leq y_i(0) \forall i$ by definition.

Assuming for the induction $\Omega(k) \leq y_i(k) \forall i, k = 1, \dots, z$

($k = z + 1$) Given Eq. (14), we can apply the inequality by Klamkin and Newman [30],

$$\prod_{i=1}^n (1 - a_i) \leq 1 - \sum_{i=1}^n a_i + \sum_{i=1}^n \sum_{j>i} a_i a_j$$

(with $a_j = \Omega(z)A_{ij}\Theta_j$) to Eq. (13), resulting in:

$$\Omega(z + 1) \leq \Omega(z) + (1 - \Omega(z)) \left[1 - \prod_{j=1}^{|V|} (1 - \Omega(z)A_{ij}\Theta_j) \right]$$

For simplicity we call

$$\Psi = 1 - \prod_{j=1}^{|V|} (1 - \Omega(z)A_{ij}\Theta_j)$$

and we obtain

$$y_i(z + 1) \geq y_i(z) + (1 - y_i(z))\Psi$$

as, because of the inductive step, $\Omega(z) \leq y_i(z) \forall i, z$.

Subtracting $\Omega(z + 1)$ from $y_i(z + 1)$ we get

$$y_i(z + 1) - \Omega(z + 1) \geq (y_i(z) - \Omega(z))(1 - \Psi)$$

as $\Omega(z) \leq y_i(z)$ and $\Psi < 1$ then $(y_i(z) - \Omega(z))(1 - \Psi) \geq 0$ and finally

$$\Omega(z + 1) \leq y_i(z + 1). \quad \square$$

Theorem 1 exploits the first order approximation of the SI model on a graph G given by Eq. (8), and applies the reception-equal property granted by Eq. (2) to derive theoretical stochastic upper and lower bounds for the probability that node i is infected, i.e., it has received the information, at time k .

A node-independent bound express the probability that a generic node has received the packet regardless of its position in the network averaged on all the possible sources of the information. These bounds can also be interpreted as bounds on the information delay expectation for each node when there is no knowledge on the information source position, or in the SI terminology, when the initial probability of infection is $y_i(0) = \frac{1}{|V|} \forall i$.

The importance of these bounds is that they are topology independent, thus they give a very powerful design tool to set the communication cycle duration and other network tuning parameters when some constraints on information dissemination should be met with high probability.

3.1. Solving the bounds

Eq. (9) is a second order difference equation similar to the *logistic map*, but its parameters keep it in the stability region (it is not chaotic), furthermore we are only interested in studying its value for $\omega(k) \in [0, 1]$. Eq. (9) has two fixed points, $\omega(k) = \{0, 1\}$. The first one is irrelevant as $\omega(0) > 0$ and the latter is an attractor as Eq. (9) is non-decreasing.

Let $\omega(k) = 1 - \epsilon_k$ be the probability that node i is infected at time k , with $\epsilon_0 = 1 - \frac{1}{|V|}$, then we have

$$\begin{aligned} \omega(k + 1) &= 1 - \epsilon_{k+1} = 2(1 - \epsilon_k) - (1 - \epsilon_k)^2 = \\ &= 2 - 2\epsilon_k - 1 + 2\epsilon_k - \epsilon_k^2 = 1 - \epsilon_k^2 \end{aligned}$$

and consequently $\epsilon_{k+1} = \epsilon_k^2$ that finally implies

$$\omega(k) = 1 - \epsilon_0^{2^k} = 1 - \left(1 - \frac{1}{|V|} \right)^{2^k} \quad (15)$$

and solving for k

$$k = \left\lceil \log_2 \left(\log_{\left(1 - \frac{1}{|V|}\right)} (1 - p) \right) \right\rceil, \quad \forall p \in \left(\frac{1}{|V|}, 1 \right) \quad (16)$$

where k is the average number of time cycles needed for a node to have received a packet with probability p .

Eq. (15) indicates that the reception-equal condition grants, regardless of the network topology of G , a double exponential speed of convergence (much faster than the exponential speed in the SI model) in the early distribution phase (when $y_i(k) \ll 1, \forall i$).

Eq. (10) is strictly non-decreasing for $\Omega(k) \in [0, 1]$ and with a slower growth than Eq. (9). Unfortunately, Eq. (10) cannot be stated in closed form, but we can numerically integrate the difference equation.

3.2. Energy consumption with RE strategy

Energy consumption is one of the key performance metrics in battery-powered networks. The RE property we are exploiting in this work (see [16]) grants that each node j sends at most $k\Theta_j$ during k cycles; furthermore it is also granted that $\Theta_j \leq |N_j|$ where N_j is the neighbor set of node j . These bounds allow a node to tune its own energy consumption: if a node is low on battery, it can simply reduce its neighbor set (dropping some links, or simply avoiding communicating their presence) and let the system recompute the optimal parameters. In the column-uniform scenario, parameter computation can be done locally gossiping the neighbor set size with neighbors [16], hence their update is energetically cheap. The actual energy consumption is also function of the packet scheduling efficiency (avoidance of duplicates) and it is outside the scope of this work.

3.3. Model limitations

The model we described is very powerful but it is not truly universal, as it relies on a set of assumptions.

The first limitation lies in the approximation introduced by Eq. (7), and can be easily explained with an example. Consider a linear network, in which every node has exactly two neighbors (excluding the nodes at the extremes) and the diameter of the network is $|V| - 1$. Then at most two nodes get infected per time cycle. In this corner case the delay needed to achieve $p = 0.9999\%$ grows linearly with $|V|$ and breaks the theoretical bounds we formulated. Eq. (7) in fact assumes that $P\{i \in S\}$ is independent of $P\{j \in I\}$ and multiplies the two probabilities. This assumption is strictly true only in a full mesh, as all nodes are neighbors so the probability of j to be infected at time k depends only on the total number of infected nodes at time $k - 1$ and not on their position. In general this is not true, as $P\{i \in S\}$ depends on $P\{j \in S\}$ if j and i are neighbors, while the inference between $P\{j \in S\}$ and $P\{i \in S\}$ decreases with the distance from j to i . In terms of density and path diversity the linear network is at the opposite extreme of a full mesh, thus it is not surprising that our model fails to capture its behavior.

The second limitation lies in the fact that our bound gives an average delay computed over all the possible sources and destinations, but it does not describe the distribution of the delay. In graphs that show some regularity (like an Erdős-Rényi graph) we expect the deviation from the mean to be narrow, but in general this may not be true.

Finally, the analysis models a zero-signaling dissemination, without any heuristic to limit duplicates or a protocol to coordinate nodes. In real applications simple heuristics can be used to improve the information diffusion (e.g., do not send the packet twice to the same neighbor, and do not send it back to the neighbor that sent it to you). Also, what is the impact of packet losses on RE strategy, do they hamper it more than traditional SE ones or vice-versa the gain achieved is even larger?

Three interesting questions arises from these model approximations and considerations:

(Q1) Do the bounds still hold when we apply the RE strategy to networks that have a density much lower than a full mesh, and hence the simple SI independence assumptions do not hold?

(Q2) How sensitive is delay to the position of the source?

(Q3) Is the RE still optimal if we introduce scheduling heuristics or consider lossy links?

Section 4 analyzes and gives insight in these three important questions.

4. Numerical results

To investigate the three questions posed above, we implemented an ad-hoc, event-driven simulator in Python. The simulator actually generates and spread packets in networks with arbitrary topology. The packet transmission time is considered much smaller than the duty-cycle as reasonable in low duty-cycle networks, and the propagation delay is negligible. In every simulation experiment one node, called source, generates one information message and sends it to a neighbor at random. From that moment on, the source behaves as any other node in the flooding process; the experiment ends at time k when all the nodes have received the message. Each experiment is repeated 20 times to account for the randomness and variability of the flooding process computing k average and variability (standard deviation).

We compare this quantity with $\omega^{-1}(p)$ and $\Omega^{-1}(p)$ for $p = 0.9999$ (assumed as certainty of reception as the bounds are stochastic and go to ∞ for $p = 1$). Bounds are normally identifies with marks, while solid lines identify RE, the reception-equal optimized strategy of Section 2.1, and SE the sending-equal standard strategy.

Furthermore, where it is meaningful to gain insight in the problem, we also report two additional “limits” or bounds. The first one is the well known $\log_2(|V|) + 1$ cycles, which defines the minimum possible delay diffusion when nodes transmit one message per cycle, and stems from the simple observation that in this case the number of message copies in the network can at most double at every cycle. The second one is the maximum distance from the source, which defines the absolute minimum diffusion time given a source and the minimum spanning tree that reaches all nodes from the source. Notice that in this case nodes are not constrained to transmit one message per cycle, but each node sends as many copies as its outgoing links in the tree.

Performance comparison with other optimization or heuristic-based techniques such as the tree-based solution proposed in [7] are outside the scope of this work, as they would require a full implementation and also to consider details and constraints that are not coherent with the generality of the theoretical approach adopted here. Nevertheless, a full implementation of a flooding protocol based on the RE strategy and its comparison with other techniques that use the same amount of transmission resources is feasible, and it is of the utmost interest as future work.

The first set of results, reported in Section 4.1, answers question Q1. To achieve this goal we use Barabási–Albert and Erdős–Rényi networks for their well-known properties; the first one is a class of preferential-attachment graphs, for which there are few nodes well connected within each other and the rest of the network, while the remaining nodes have few peripheral connections (this is the class of a large part of real-world networks, like the internet and social networks). Barabási–Albert graph generators take in input a parameter m indicating for each node how many outgoing link are setup. The Erdős–Rényi type of network is a model for which links are independently randomly picked with a probability p ; this class has been widely studied as it allows an easy derivation of statistics (this is typically the class of P2P overlay networks). We vary the network parameters, and we show that our theoretical bounds hold even in networks with density far from a full mesh. The second set of results, shown in Section 4.2, answers question Q2. We show that our bounds are still valid even when we consider the results for each single source. Finally, the third set of results, in Section 4.3, answers question Q3. We show that packet losses do not hamper the bounds, while simple heuristics applied to message scheduling improve RE-based results bringing them, as expected, below the scheduling agnostic stochastic threshold. As a further comparison and validation we introduce, in Section 4.4, a third family of graphs (Waxman) with completely different characteristics. Waxman networks are created placing nodes randomly in a rectangular area and then placing links between them with a probability p_w exponentially dependent on their distance d : $p_w = \alpha \cdot e^{-\frac{d}{\beta}}$, where α, β are free parameters.

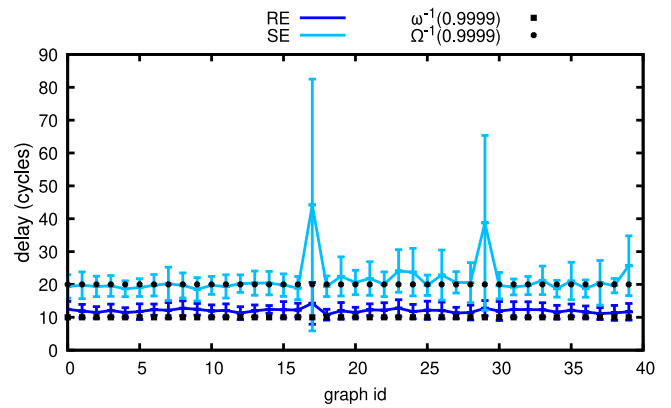


Fig. 2. Flooding delay mean and standard deviation on Barabási–Albert networks of 100 nodes.

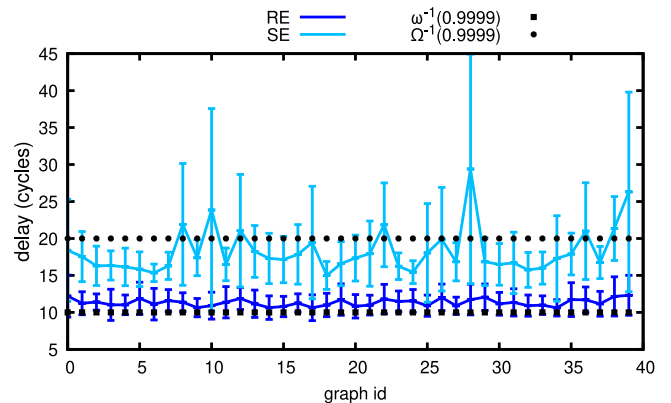


Fig. 3. Flooding delay mean and standard deviation on Erdős–Rényi networks of 100 nodes.

4.1. Results on low density graphs

We simulate Erdős–Rényi and Barabási–Albert networks with a low and approximately constant edge density $\frac{|E|}{|V|} \simeq 4$. For each network type, we execute 100 different experiments (each consisting of 20 repetitions) picking a different node as the source, thus each experiment (one point in a graph) consists of 2000 different simulations.

We first show, in Figs. 2 and 3, the flooding performance across multiple graph realizations using the same graph parameters. The x axis consists of 40 different network realizations (graph id). The figures show that the point estimate of the average time k to complete the flooding with RE strategy falls between the bounds, which confirms that the theoretical bounds hold beyond the approximations needed to obtain them. On the other hand SE-based dissemination performs poorly, constantly around the upper bound (i.e., larger distribution time Ω^{-1}) of RE. Moreover, SE strategy, failing to adapt to topological properties of the network, sometimes (e.g., graphs 17 and 29 in Fig. 2) display an extremely poor performance while RE adapts to the graph characteristics maintaining roughly constant performance.

Figs. 4 and 5 explore the influence of the network size on the RE and SE strategies and on the bounds on RE. We explore Barabási–Albert and Erdős–Rényi networks with up to 2000 nodes; the bounds hold independently of the network size for the point estimate and for the standard deviation. On the other hand SE-based dissemination performance degrades with the network size and significantly deviates even from the lower performance bound (Ω) of RE.

Figs. 4 and 5 report also another bound (see [31] for a distributed scheduling achieving it for streaming applications) that defines the lowest possible delay in a full-mesh network with complete knowledge

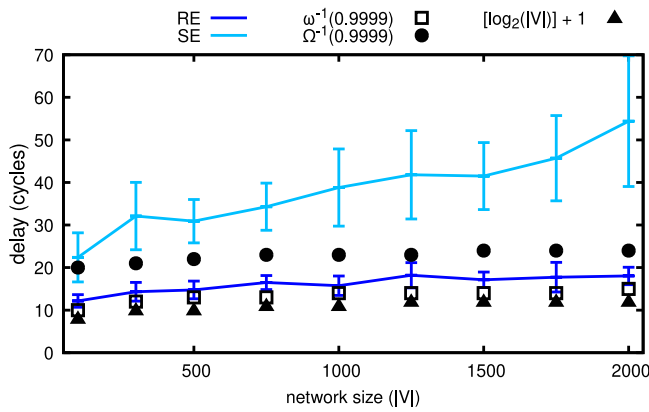


Fig. 4. Mean and standard deviation (as barbs) for the dissemination delay and theoretical bounds for Barabási–Albert networks varying $|V|$.

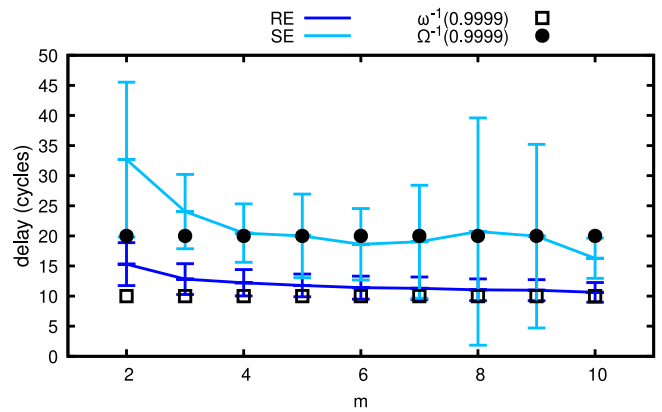


Fig. 6. Flooding delay mean and standard deviation on Barabási–Albert networks of 100 nodes varying the “m” parameter of the distribution.

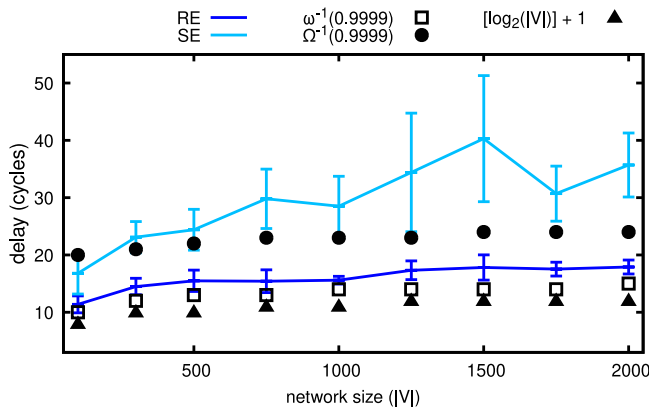


Fig. 5. Mean and standard deviation (as barbs) for the dissemination delay and theoretical bounds for Erdős–Rényi networks varying $|V|$.

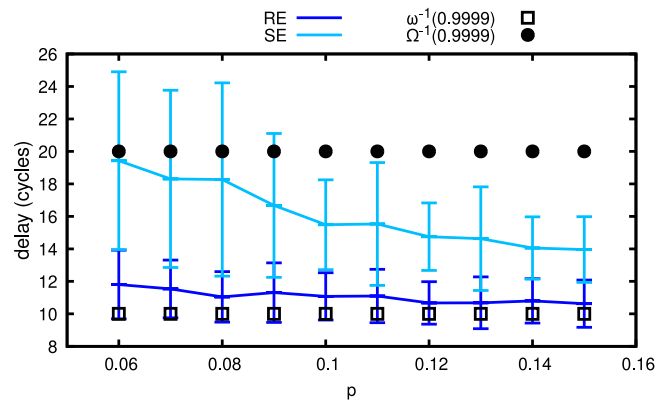


Fig. 7. Flooding delay mean and standard deviation on Erdős–Rényi networks of 100 nodes varying the distribution “p” parameter.

of nodes state (each node knows what are the nodes that have not yet received the message) for SE strategies. Such bound is important to our study as we can compare state-agnostic RE strategies and bounds with a bound that requires full-knowledge of the network state to be achieved. It is interesting how ω^{-1} is close to this bound, which suggests that an optimal scheduling associated with an RE strategy can actually improve the flooding delay beyond the minimum achievable with SE strategies without using (globally) more resources. This would actually move the lowest achievable delay bound beyond the state-of-art for unstructured flooding distribution, at the only cost of allocating transmission resources non uniformly among nodes.

Results presented so far are obtained maintaining the graph density constant to ease the comparisons. Figs. 6 and 7 present flooding delay results varying the graph distribution parameters, which produce networks of growing density. The “m” parameter in Barabási–Albert networks refers to the number of bi-directional links that every new node added to the network establishes with already existing nodes. The “p” parameter in Erdős–Rényi networks is the probability of adding a link with any of the other nodes. RE steadily performs better than SE and the point estimate always falls between the two bounds, while in some cases (specific network realization with some given source node) the distribution delay is even smaller than the ω^{-1} bound of the average, which is however fully admissible from a theoretic point of view. The impressive reduction in source dependent variability compared to SE is remarkable.

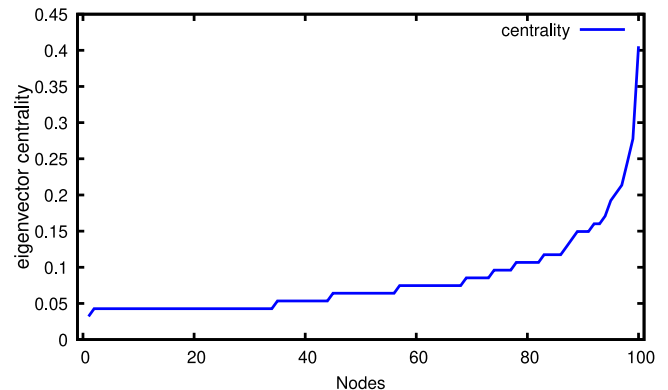


Fig. 8. Ordered eigenvector centrality of nodes in a Barabási–Albert network of 100 nodes.

4.2. Dependency on the source position

Here we answer Q2 and present insights on the correlation between flooding time and the position of the source node. Once verified that our bound holds on average, we want to show that it also gives a good indication on the flooding performance of each possible source.

We first focus on a specific graph to clarify RE properties. Fig. 8 reports the ordered values of the eigenvector centrality for the nodes of a 100-nodes Barabási–Albert graph, and Fig. 9 shows the flooding delay simulated placing the source in each node (point estimate of the average and standard deviation barbs) with the nodes in the same

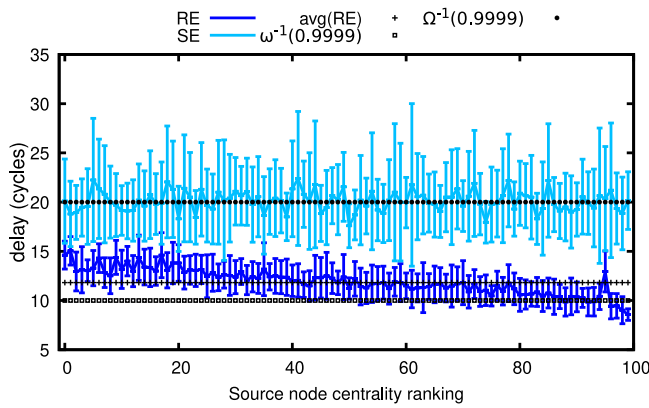


Fig. 9. Flooding delay mean and standard deviation on a Barabási–Albert network of 100 nodes varying the source node centrality (source nodes are sorted with respect their centrality value). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

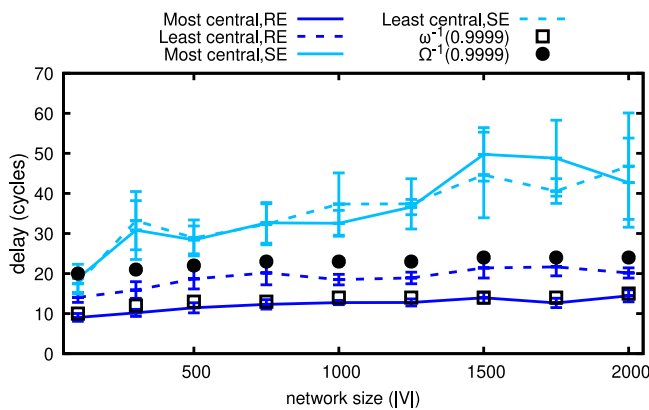


Fig. 10. Flooding delay mean and standard deviation on Barabási–Albert networks using the most and least central nodes as sources for RE and SE strategies.

order of Fig. 8. In contrast to the SE strategy, the RE strategy flooding delay is sensitive to source node centrality (compare the blue curve and marks with the crosses that indicate the average value), but the per-node average performance always falls between the bounds, excluding a very few central nodes for which the performance is better than the average lower bound. It is also worth noticing how, regardless of node centrality, RE flooding delay performances are clearly separated from SE values, excluding the possibility of casual results and highlighting the fact that the gain in average is not achieved penalizing less central nodes.

To further extend the validation, Figs. 10 and 11 report the flooding delay on Barabási–Albert and Erdős–Rényi networks of growing size with constant density ($\frac{|E|}{|V|} \simeq 4$). For each graph we show the performance of the most central and the least central nodes, and we compare RE with SE and report the RE bounds as well. These results confirm the previous findings: RE performance falls between the bounds in almost all the cases; the few exceptions in Barabási–Albert graphs perform better than the lower bound, while in Erdős–Rényi graphs the least central sources slightly exceed the upper average bound. For any size RE performs significantly better than SE, which appears to be less dependent on the centrality of the source, but it is always worse than the upper bound on the RE strategy.

4.3. Results with scheduling heuristics

Results in the previous sections show how the RE strategy lies within the stochastic bounds in generic topologies and always outperforms

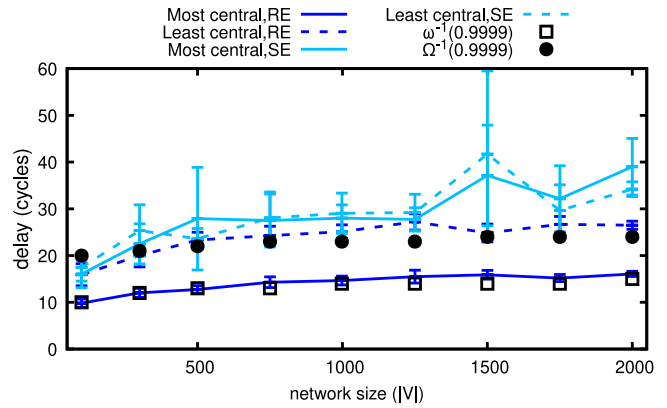


Fig. 11. Flooding delay mean and standard deviation on Erdős–Rényi networks using the most and least central nodes as sources for RE and SE strategies.

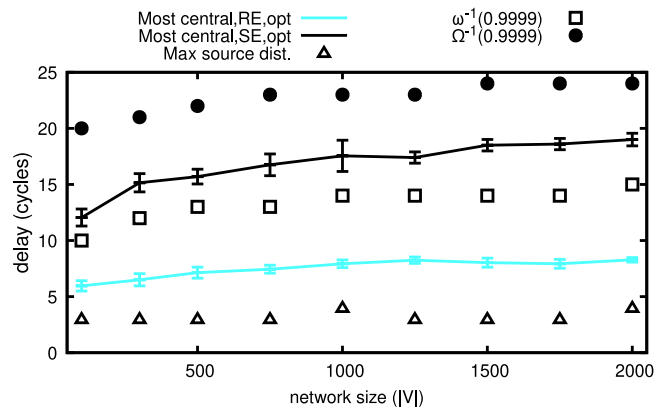


Fig. 12. Flooding delay mean and standard deviation on Barabási–Albert networks when the source is the most central nodes and scheduling heuristics are applied.

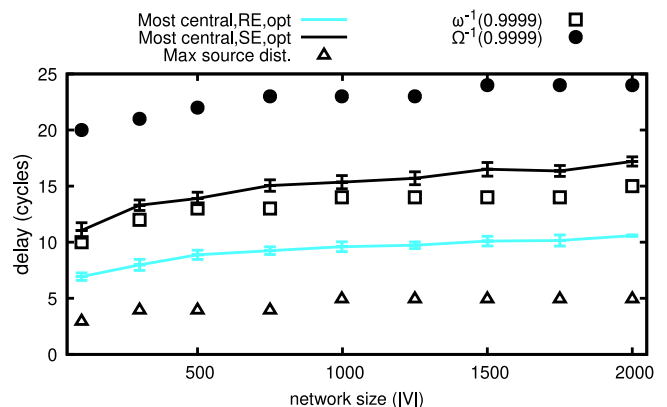


Fig. 13. Flooding delay mean and standard deviation on Erdős–Rényi networks when the source is the most central nodes and scheduling heuristics are applied.

SE. As our model cannot capture some real world effects, we still have to answer question Q3. In this section we introduce some simple scheduling heuristics and show that the RE strategy is further improved, in other words, that any intelligent flooding protocol will likely benefit from the adoption of RE.

Figs. 12 and 13 report performance result comparison between SE and RE strategies using simple scheduling heuristic: *do not send the packet twice to the same neighbor, and do not send it back to the neighbor that sent it to you*. The figures refer to Barabási–Albert and Erdős–Rényi graphs of growing size, and are similar to Figs. 4 and 5, but they show

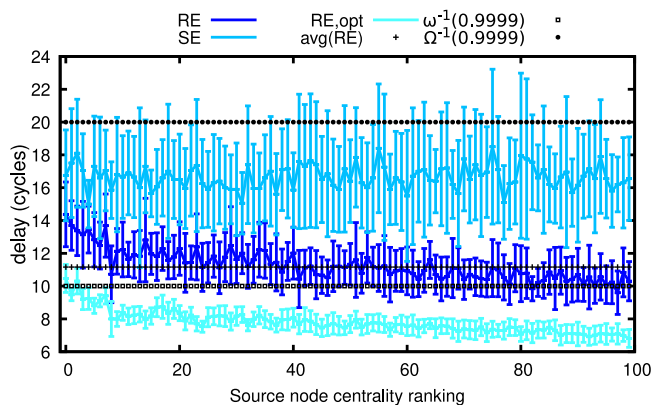


Fig. 14. Flooding delay mean and standard deviation on Waxman network of 100 nodes as a function of the source node centrality (source nodes are sorted with respect their centrality value).

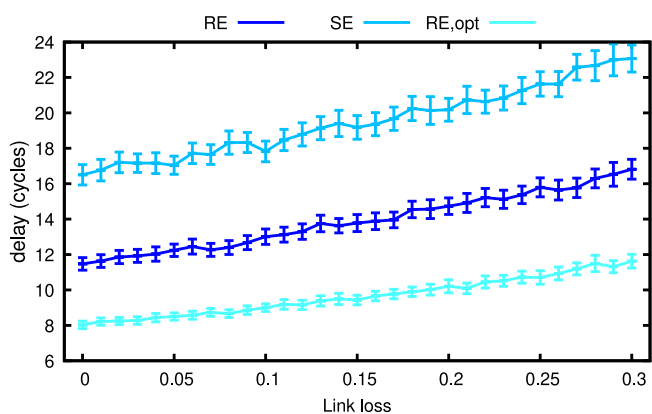


Fig. 15. Flooding delay confidence interval (at 99%) on a Waxman network of 100 nodes with a Bernoulli link loss model.

only the behavior of the most central node in the network. This makes it possible to compare RE performance with another meaningful bound: the maximum distance from the source. In practice, we are comparing RE with the delay achievable flooding the information on a minimum depth tree rooted in the source. Note that to achieve this bound, in practice, the network should be able to build a tree rooted on the source and maintain it at every topology change, which is a very costly operation in a network with frequent topology changes. Yet RE, which is fully distributed and needs only local interactions for maintenance, performs surprisingly close to this optimum in both kinds of graphs.

4.4. Results on Waxman graphs with losses

As a further validation step we simulate a different family of graphs: Waxman graphs. Barabási–Albert and Erdős–Rényi graphs do not support any notion of space (and hence are easy to generate and manipulate). Nodes do not have spatial coordinates, and thus, they are all treated equally. Real world networks are instead spatial networks, and their characteristics differ from synthetic graphs, for instance, they introduce border effects [32]. Among the many models available for spatial networks, Waxman graphs are the simplest model used to represent communication networks [33]. For our tests we use Waxman graphs generated with parameters $\alpha = 0.4, \beta = 0.2$, granting the density $\frac{|E|}{|V|} \simeq 4$.

Fig. 14 reports the same results presented in Fig. 9 obtained on a Waxman graph of 100 nodes. The results follow the same trend of Fig. 9 with a slightly more noisy behavior, which confirms the bounds hold also on graph whose characteristics are far from the infection model.

We also report the performance of RE with scheduling heuristics, which confirms a clear improvement.

As a final step, we also explore the performance with packet losses. Given a Waxman graph, we simulate flooding with a loss probability on the links that follows a Bernoulli distribution with average link loss p_l and we measure the propagation delay. Fig. 15 reports the flooding delay confidence intervals with 99% confidence level considering a link loss probability $p_l \in [0, 0.3]$. As expected at this point, even with extremely high values of p_l , RE performs constantly better than SE and applying scheduling heuristics further improve it, also reducing the standard deviation.

5. Related works

Low-duty-cycle WSNs are said *synchronous* if the node active state happens at a fixed time or *asynchronous* if they are scheduled independently. Literature focuses mostly on asynchronous low-duty-cycle WSNs, where often tree-based dissemination or collection overlays are built. This overview of literature focuses both on recent papers that address the problem of information flooding (not collection to a sink) in WSNs, and on papers that analyze or propose epidemic dissemination techniques, in this case not limited to WSNs, but ranging also from P2P networks to classical infection models used in, or derived from, medical literature.

5.1. Flooding in WSNs

Wang and Liu [14] propose a reinterpretation of flooding for the context of WSNs and they provide a centralized optimization model from which they derive an approximated distributed solution. Flooding has been also investigated by Cao et al. using Fountain coding [34]. Our approach is completely distributed.

While RE works with unstructured mesh networks, several works have been proposed for flooding on tree networks; Guo et al. [7] address both the delay and the energy constraints deriving a tree-based distribution solution considering lossy links. Cheng et al. [19] propose a flooding tree construction algorithm optimized with respect to the energy consumption, but also considering delay bounds. This algorithm is an approximated distributed version of a centralized optimal one. The work by Niu et al. [35] follows the same scheme as they propose a heuristic algorithm derived from a minimum spanning tree centralized model. Yan et al. [36] investigate the potential of network coding in the context flooding using trees.

There are also works optimizing existing flooding solutions; Cheng et al. [6] propose the Dynamic Switching-based Reliable Flooding (DSRF) to enhance the reliability of flooding. The flooding optimization by Guo et al. [37] synchronize the active state of nodes sharing the same parent node in a tree. Physical channel overhearing has been investigated by Xu et al. [38] as a mean to save delay during message flooding. These works can be used on top of other strategies.

Asynchronous Duty-cycle Broadcasting (ADB) is a protocol implemented directly in the MAC layer of WSN nodes which allows flooding by exploiting MAC-layer information. This however implies the tied coupling of the cross-layer approaches.

In contrast to recent publications on flooding on WSN, our approach is fully decentralized and works with unstructured mesh networks without the aid of trees. That grants a higher degree of robustness against node failure, a lower signaling overhead, and promising applications in time-varying networks. One limitation to be explored in future work is the impact and exploitation of broadcast communications on wireless channels.

5.2. Epidemics and networks

Epidemics, the field about modeling and analyzing the dynamics of virus spreading, has been prolific in the past decades, though only recently proper insights on how to control it has been provided [26]. A large part of computer science literature on epidemics focus on malware spreading [27,28,39]. Chen et al. [28] use the Susceptible–Infected–Recovered (SIR) model to control dissemination of information in heterogeneous, time-varying networks. The work by Dadlani et al. [27] uses a SIS model instead and provide infection stability results. This work, together with the one by Ganesh et al. [40] highlight the importance of being dependent on a specific network topology for studying epidemics. That is a crucial observation that our approach overcomes exploiting the reception-equal property obtained with the re-assignment of resources based on the eigenvector centrality.

Works by Liu and Buss [41] and by Ogura and Preciado [42] use the SIS model for data dissemination; the former optimizing the node transmission rate while the latter defining exponential growth conditions for time-varying networks. Chen et al. [43] uses epidemics to model and analyze information spreading in sensor networks.

Other papers deal with different aspects of data dissemination through epidemics approaches. The paper by Chen et al. [29] focuses on delivery dynamics on WSN with cognitive radios, the work by Ramanathan et al. [44] optimize the loss rate for Delay Tolerant Network (DTN) and Byun and So [45] address the context of duty-cycled Wireless Sensor-Actuator Network (WSAN) and propose a scheme to adjust the node transmission rates for user-given delay constraint.

None of the aforementioned papers, in part also because of their application fields, give delays bounds on information flooding that are independent of the network topology. Up to now it was considered that the optimal strategy to flood information to all nodes of a network could not be independent of its topology. The results we are presenting, instead, show that it is possible to exploit the topological properties of the network to decorrelate the optimal flooding strategy from the topology itself. This observation is what enables the general analysis that in this paper leads to the bounds presented in Section 3.

6. Conclusions and further work

Flooding information to all the nodes of a network remains an important function in many networks and applications. Many solutions have been proposed and are working satisfactorily in networks from P2Ps overlays to WSNs, but in many cases they require a non-marginal overhead to build a distribution tree, or they are fragile to topology changes. This paper presented fundamental delay bounds for epidemic flooding in low duty-cycle networks that exploit the eigenvector centrality of nodes in the network to allocate resources, i.e., how many copies of the information per time-cycle a node must send, and to whom of its neighbors. The bounds apply to a resource allocation strategy that we have called *reception equal* (RE), and they are independent of the network topology. They show that with RE flooding the lower bound on delay converges with double exponential speed, while the upper bound is exponentially fast, thus ensuring that a proper protocol designed on these properties will converge at least exponentially fast. Furthermore, the results are constructive, i.e., they indicate a path to realize a protocol that obtains a performance within the bounds.

Theoretical bounds on complex graph structures are in general very difficult to derive, and in fact our bounds are valid on average for all sources and destinations in the graph. To move a step forward towards a real application we have implemented the RE flooding in a simulator and showed the bounds hold with a very good approximation also to estimate the performance of flooding from a single source. We discussed simple heuristics that improve on the bound, exploiting additional knowledge that was not included in the theoretical model for mathematical tractability, and we have tested the strategy with a simple loss model. In all these settings, our results are extremely encouraging.

Future work on this study include the use of RE flooding for specific applications like time synchronization in sensor networks, a comparison with state-of-the-art protocols based on trees or other dissemination structures, and further theoretical analysis of the optimized protocol.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] L. Baldesi, L. Maccari, R. Lo Cigno, On the properties of infective flooding in low-duty-cycle networks, in: 15th IFIP/IEEE Conf. on Wireless on Demand Network Systems and Services, WONS'19, Wengen, CH, 2019, pp. 1–8.
- [2] P. Aggarwal, H. Gossain, A. Joshi, M. Shukla, System and method for selecting a route based on link metrics incorporating channel bandwidth, spatial streams and/or guard interval in a multiple-input multiple-output (MIMO) network, 2014, Google Patents.
- [3] L. Baldesi, L. Maccari, R. Lo Cigno, Optimized cooperative streaming in wireless mesh networks, in: IFIP TC-6 Int. Networking Conf., IFIP NETWORKING'16, Vienna, Austria, 2016, pp. 350–358.
- [4] Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, S. Shenker, Making gnutella-like P2P systems scalable, in: Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM'03, 2003, pp. 407–418.
- [5] N. Magharei, R. Rejaie, Y. Guo, Mesh or multiple-tree: A comparative study of live P2P streaming approaches, in: 26th IEEE Int. Conf. on Computer Communications, INFOCOM'07, Barcelona, Spain, 2007, pp. 1424–1432.
- [6] L. Cheng, J. Niu, Y. Gu, C. Luo, T. He, Achieving efficient reliable flooding in low-duty-cycle wireless sensor networks, *IEEE/ACM Trans. Netw.* 24 (6) (2016) 3676–3689.
- [7] S. Guo, L. He, Y. Gu, B. Jiang, T. He, Opportunistic flooding in Low-Duty-Cycle wireless sensor networks with unreliable links, *IEEE Trans. Comput.* 63 (11) (2014) 2787–2802.
- [8] K. Han, J. Luo, Y. Liu, A.V. Vasilakos, Algorithm design for data communications in duty-cycled wireless sensor networks: A survey, *IEEE Commun. Mag.* 51 (7) (2013) 107–113.
- [9] L. Atzori, A. Iera, G. Morabito, The internet of things: A survey, *Comput. Netw.* 54 (15) (2010) 2787–2805, Elsevier.
- [10] L. Xu, W. He, S. Li, Internet of things in industries: A survey, *IEEE Trans. Ind. Inform.* 10 (4) (2014) 2233–2243.
- [11] Z. Li, M. Li, J. Liu, S. Tang, Understanding the flooding in Low-Duty-Cycle wireless sensor networks, in: 40th IEEE Int. Conf. on Parallel Processing, ICPP, Taipei City, Taiwan, 2011, pp. 673–682.
- [12] H.-C. Hsiao, C.-P. He, A tree-based peer-to-peer network with quality guarantees, *IEEE Trans. Parallel Distrib. Syst.* 19 (2008) 1099–1110.
- [13] D. Carra, R. Lo Cigno, E.W. Biersack, Stochastic graph processes for performance evaluation of content delivery applications in overlay networks, *IEEE Trans. Parallel Distrib. Syst.* 19 (2008) 247–261.
- [14] F. Wang, J. Liu, On reliable broadcast in low duty-cycle wireless sensor networks, *IEEE Trans. Mobile Comput.* 11 (5) (2012) 767–779.
- [15] R. Lo Cigno, M. Nardelli, M. Welzl, SESAM: A Semi-synchronous, energy savvy, application-aware MAC, in: 6th IFIP/IEEE Int. Conf. on Wireless on-Demand Network Systems and Services, WONS'09, 2009, pp. 93–100.
- [16] L. Baldesi, L. Maccari, R. Lo Cigno, On the use of eigenvector centrality for cooperative streaming, *IEEE Commun. Lett.* 21 (9) (2017) 1953–1956.
- [17] M. Newman, *Networks: An Introduction*, Oxford University Press New York, 2010.
- [18] M. Ajelli, R. Lo Cigno, A. Montresor, Modeling botnets and epidemic malware, in: IEEE Int. Conf. on Communications, ICC, Cape Town, South Africa, 2010, pp. 1–5.
- [19] L. Cheng, J. Niu, C. Luo, L. Shu, L. Kong, Z. Zhao, Y. Gu, Towards minimum-delay and energy-efficient flooding in low-duty-cycle wireless sensor networks, *Comput. Netw.* 134 (2018) 66–77.
- [20] K.S. Yildirim, A. Kantarci, Time synchronization based on slow-flooding in wireless sensor networks, *IEEE Trans. Parallel Distrib. Syst.* 25 (1) (2014) 244–253.
- [21] P. Dutta, D. Culler, Practical asynchronous neighbor discovery and rendezvous for mobile sensing applications, in: 6th ACM Conf. on Embedded Network Sensor Systems, SenSys'08, Raleigh, NC, USA, 2008, pp. 71–84.
- [22] J. Zhang, C. Yang, X. Zhang, A high-bandwidth live streaming model in mesh-based peer-to-peer networks, *IEEE Commun. Lett.* 20 (12) (2016) 2390–2393.
- [23] P. Bonacich, Some unique properties of eigenvector centrality, *Soc. Netw.* 29 (4) (2007) 555–564, Elsevier.

- [24] L. Page, S. Brin, R. Motwani, T. Winograd, *The PageRank Citation Ranking: Bringing Order to the Web*, Tech. Rep., Stanford InfoLab, 1999.
- [25] H. Ishii, R. Tempo, Distributed randomized algorithms for the PageRank computation, *IEEE Trans. Autom. Control* 55 (9) (2010) 1987–2002.
- [26] C. Nowzari, V. Preciado, G. Pappas, Analysis and control of epidemics: A survey of spreading processes on complex networks, *IEEE Control Syst. Mag.* 36 (1) (2016) 26–46.
- [27] A. Dadlani, M. Kumar, K. Kim, K. Sohraby, Stability and immunization analysis of a malware spread model over scale-free networks, *IEEE Commun. Lett.* 18 (11) (2014) 1907–1910.
- [28] P. Chen, S. Cheng, K. Chen, Optimal control of epidemic information dissemination over networks, *IEEE Trans. Cybern.* 44 (12) (2014) 2316–2328.
- [29] P. Chen, S. Cheng, H. Hsu, Analysis of information delivery dynamics in cognitive sensor networks using epidemic models, *IEEE Internet Things J.* 5 (4) (2018) 2333–2342.
- [30] M. Klamkin, D. Newman, Extensions of the Weierstrass product inequalities, *Math. Mag.* 43 (3) (1970) 137–141.
- [31] L. Abeni, C. Kiraly, R. Lo Cigno, On the optimal scheduling of streaming applications in unstructured meshes, in: 8th IFIP-TC 6 Int. Networking Conf., NETWORKING 2009, Aachen, Germany, 2009, pp. 117–130.
- [32] M. Barthélemy, Spatial networks, *Phys. Rep.* 499 (1) (2011) 1–101.
- [33] B.M. Waxman, Routing of multipoint connections, *IEEE J. Sel. Areas Commun.* 6 (9) (1988) 1617–1622.
- [34] Z. Cao, J. Wang, D. Liu, X. Miao, Q. Ma, X. Mao, Chase++: Fountain-enabled fast flooding in asynchronous duty cycle networks, in: IEEE Conference on Computer Communications, INFOCOM'18, Honolulu, HI, USA, 2018, pp. 2276–2284.
- [35] J. Niu, L. Cheng, Y. Gu, J. Jun, Q. Zhang, Minimum-delay and energy-efficient flooding tree in asynchronous low-duty-cycle wireless sensor networks, in: IEEE Wireless Communications and Networking Conference, WCNC, Shanghai, China, 2013, pp. 1261–1266.
- [36] F. Yan, X. Zhang, H. Zhang, Flooding with network coding under a schedule-based spanning tree in low-duty-cycle wireless sensor networks, *IEEE Wirel. Commun. Lett.* 7 (2) (2018) 270–273.
- [37] S. Guo, S.M. Kim, T. Zhu, Y. Gu, T. He, Correlated flooding in low-duty-cycle wireless sensor networks, in: 19th IEEE Int. Conf. on Network Protocols, ICNP, Vancouver, BC, Canada, 2011, pp. 383–392.
- [38] L. Xu, G. Chen, J. Cao, S. Lin, H. Dai, X. Wu, F. Wu, Optimizing energy efficiency for minimum latency broadcast in low-duty-cycle sensor networks, *ACM Trans. Sensor Netw.* 11 (4) (2015) 57.
- [39] M. Garetto, W. Gong, D. Towsley, Modeling malware spreading dynamics, in: 22nd IEEE Int. Conf. on Computer and Communications, INFOCOM'03, San Francisco, CA, USA, 2003, pp. 1869–1879.
- [40] A. Ganesh, L. Massoulié, D. Towsley, The effect of network topology on the spread of epidemics, in: 24th IEEE Int. Conf. on Computer and Communications, INFOCOM'05, Miami, FL, USA, 2005, pp. 1455–1466.
- [41] F. Liu, M. Buss, Optimal control for information diffusion over heterogeneous networks, in: 55th IEEE Conference on Decision and Control, CDC'16, Las Vegas, NV, USA, 2016, pp. 141–146.
- [42] M. Ogura, V. Preciado, Stability of spreading processes over time-varying large-scale networks, *IEEE Trans. Netw. Sci. Eng.* 3 (1) (2016) 44–57.
- [43] P. Chen, S. Cheng, H. Hsu, Analysis of information delivery dynamics in cognitive sensor networks using epidemic models, *IEEE Internet Things J.* 5 (4) (2018) 2333–2342.
- [44] R. Ramanathan, R. Hansen, P. Basu, R. Rosales-Hain, R. Krishnan, Prioritized epidemic routing for opportunistic networks, in: 1st Int. MobiSys Workshop on Mobile Opportunistic Networking, MobiOpp'07, San Juan, Puerto Rico, 2007, pp. 62–66.
- [45] H. Byun, J. So, Node scheduling control inspired by epidemic theory for data dissemination in wireless sensor-actuator networks with delay constraints, *IEEE Trans. Wirel. Commun.* 15 (3) (2016) 1794–1807.