

Broadcast Flooding Revisited: Survivability and Latency

Petteri Mannersalo
VTT Technical Research Centre of Finland
Email: petteri.mannersalo@vtt.fi

Alireza Keshavarz-Haddad
Department of Electrical and
Computer Engineering
Rice University
Email: alireza@rice.edu

Rudolf Riedi
Department of Statistics
Rice University
Email: riedi@rice.edu

Abstract—This paper addresses the dynamics of broadcast flooding in random wireless ad hoc networks. In particular, we study the subset of nodes covered by a flood as well as timing issues related to the first (latency) and the last time (duration of back-chatter) at which a broadcast is received by a fixed node. Notably, this analysis takes into account the MAC-layer as well as background traffic which both are often neglected in related studies.

Assuming a protocol model for the transmission channel which accounts for carrier sensing and interference, we find bounds for the probability of survival of the flood and for its coverage probabilities. Moreover, under certain conditions on the parameters, we establish asymptotical linear bounds on the latency as the distance from the origin of the flood increases and show that the duration of the back-chatter is stochastically bounded. The analytical results are compared to simulation.

I. INTRODUCTION

Broadcast is one of the basic functionalities in mobile wireless Ad Hoc networks. Many routing protocols employ broadcast for establishing new routes in mobile environment; examples include Dynamic Source Routing (DSR), Ad hoc On-demand Distance Vector (AODV), Zone Routing Protocol (ZRP), and Location Aided Routing (LAR) [1]. In addition, it is commonly used in network self-configuration and data gathering. In a mobile setting, updates based on broadcast should take place frequently to keep the network functioning. In some alarm networks, all point-to-point communications may be based on broadcast.

The most simple broadcast mechanism is flooding where any node receiving the broadcast packet forwards it to all its neighbors. Flooding is known to work fine in sparse and moderately dense networks. On the other hand, redundancy, collisions and contention constitute serious issues in highly dense networks, first addressed under the name *broadcast storm problem* by Ni et al. [2]. Several schemes have been proposed to avoid the broadcast storm problem. These usually include building a broadcast backbone, i.e., only a subset of the nodes forward the packet (see [3], [4] for references).

If not operating in an extremely dense network, one may ask if the optimization to decrease redundancy, collisions and contention really pays off. In very dynamic mobile environments, the effort to build a “nearly optimal” backbone has to be invested repeatedly, which adds to overhead and latency,

whereas by flooding a broadcast packet would spread approximately as fast and sometimes even more reliably [3], [5]. Moreover, it should also be noticed that some of the backbone based schemes require that the node location information is first gathered by flooding.

The packets of a flood, i.e., the broadcast packets are treated differently than unicast packets in the MAC layer. A broadcast packet is usually sent to all neighbors simultaneously. Under the IEEE MAC 802.11 Distributed Coordination Function (DCF) [6], broadcast transmissions are implemented so that the broadcast packets are sent to the neighbors as soon as the radio channel is sensed to be free (carrier sensing). However, no collision detection is used to guarantee successful receipt of the packets at the destination nodes. Some studies have suggested reliable MAC designs for broadcast transmission [7], [8], [9], [10]. These schemes, however, add a large overhead to the MAC layer and have not been applied in the current standards.

We have found only one previous work analytically studying flooding in ad hoc networks. Viswanath and Obraczka [11] model the MAC layer by considering a constant probability for successful transmissions and analyze the survivability of flooding by applying Markov chain techniques. Unfortunately, their method does not capture well the geometric properties of flooding.

The main goal of this paper is to analyze behavior of broadcast flooding in sparse and moderately dense networks under light to moderate background traffic. According to our knowledge, this is the first time when MAC and background traffic are taken into account in an analytical study of flooding. The analysis is based on percolation theory, an approach pioneered in the context of Ad Hoc networks by Dousse, Franceschetti, Thiran, et al. (see [12], [13], [14]).

The contributions of the paper are the following:

- We introduce a new communication model for wireless Ad Hoc networks. The model takes into account carrier sensing and interference caused by background traffic;
- We derive bounds for the survivability and coverage probabilities. The bounds depend on the underlying percolation functions which can be numerically approximated;
- We derive characterizations of latency and back-chatter, i.e., the times to first and last reception of the broadcast

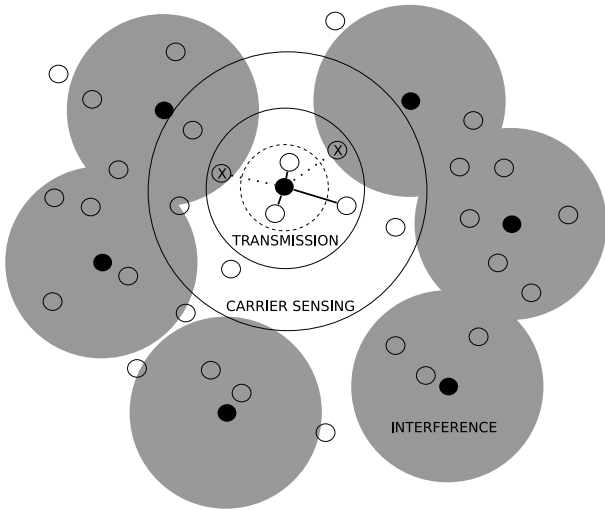


Fig. 1. Illustration of the communication model. Transmitting nodes are represented by filled, black circles, all other nodes, receiving or idle, by circles (not filled). An “X” inside a circle indicates a collision at that receiving node. For the sending node close to the middle, the circles of radii $R_C > R_T > R_F$ (R_F dashed) are drawn as well as dashed lines for failing and solid lines for successful transmissions. For the other transmitting nodes gray disks of radius R_I are indicated.

at a given node. Latency behaves linearly as the distance to the origin of the broadcast increases and back-chatter is stochastically bounded by a finite random variable.

- The analytical results are compared to ns-2 simulations.

The remainder of the paper is organized as follows. The channel model and the overall setting is described in Section II. In Section III, bounds for survivability and coverage probabilities are derived. Latency and back-chatter are analyzed in Section IV. In Section V, ns2 simulations are presented. The paper is concluded in Section VI.

II. SETTING

In the following we set the stage for our analysis: we explain the models used for network, background traffic and broadcast, and recall the basics of continuum percolation.

A. Network Model

Throughout the paper we consider an infinite planar network where the locations of the nodes are defined by a Poisson point process in \mathbf{R}^2 with a homogeneous finite *intensity* λ . By this we mean that the number of nodes lying in a measurable set A with area $|A|$ is a Poisson random variable with mean $\lambda \cdot |A|$ and is independent of the node locations outside A .

B. Communication model

We consider a communication model parameterized by *transmission range* R_T , *carrier sensing range* R_C and *interference range* R_I . The model imposes that at any time, for any sending node S and any non-sending node D the following conditions hold:

- M1 (Carrier sensing) All simultaneously transmitting nodes are at least at a distance R_C from S .

- M2 (Out of range) If D is at a distance larger than R_T from S then it can not receive the transmission from S .

- M3 (Interference) If there is a third node within a distance R_I of D which transmits at the same time then a collision occurs and D can not receive the transmission from S .

- M4 (Successful transmission) The transmission of S is successfully received at D if both, M2 and M3 fail.

Let us elaborate now on natural constraints between the parameters. Naturally, if two nodes are close enough to communicate with each other, then they should cause interference to each other as well; similarly, when they are close enough to interfere with each other then carrier sensing should be able to indicate the interference. Therefore, we assume that $R_T < R_I < R_C$. Consequently, there is a *collision-free* disk around each node with radius

$$R_F \doteq \min \{R_T, R_C - R_I\},$$

i.e., once a node has its turn to transmit, all the nodes inside the disk of radius R_F will receive the message correctly, irrespective of any other simultaneous transmissions allowed by the MAC layer. Note, that some nodes outside the collision free disc may receive the message as well. The communication model is illustrated in Fig. 1.

Comparing with other existing channel models can further instruct a sensible choice of the parameters. The channel model based on the Signal to Noise plus Interference Ratio (SINR-model, see e.g. [15]) can be used as such a reference model. In the SINR-model, node X_i can transmit to node X_j if

$$\frac{P\ell(|X_i - X_j|)}{N_0 + \gamma \sum_{k \in I \setminus i} P\ell(|X_k - X_j|)} \geq \beta,$$

where P is the transmission power, ℓ is an attenuation function, N_0 is the ambient noise, γ is a protocol parameter, I is the set of all currently transmitting nodes and β is a given threshold. Now, the carrier sensing can be implemented in such way that node X_i may send a packet only if

$$N_0 + \gamma \sum_{k \in I \setminus i} P\ell(|X_k - X_i|) \leq \xi,$$

where $\xi > N_0$.

Notably, in the SINR-model regions of transmission, carrier sensing and interference are typically not circular. In fact they may depend on the locations of all nodes and have different shapes around each node. However, the SINR-model implies constraints which the parameters R_T and R_C must satisfy. Assuming a non-increasing attenuation ℓ , we have

$$R_T \leq \ell^{-1}(N_0\beta/P), \quad (1)$$

$$R_C \geq \ell^{-1}\left(\frac{\xi - N_0}{\gamma P}\right). \quad (2)$$

The right hand side of the inequality (1) is the maximum transmission range assuming that all nodes, except the transmitting node, are idle. The inequality (2) is based on the scenario of one transmitting node and one interfering node. The inference

range R_I depends on R_T , but without further knowledge on the node locations, it could be chosen to be anything between 0 and ∞ .

C. Broadcast communications

We assume that broadcast flooding is implemented as follows:

- Once a node has received a broadcast message for the first time, it retransmits the packet to its neighbors as soon as the radio channel is found to be free (carrier sensing). No further action is taken.
- Notably, if the same message is later heard again, it is not re-sent.
- Also, no collision detection is used. Thus, the broadcast is received exactly by the nodes within range and without interference, as specified in the channel model.

This corresponds to the implementation in IEEE 802.11.

D. Traffic and MAC layer

For some results we require the following assumptions on the MAC layer and the network traffic (i.e., broadcast and background traffic together):

- A1 (Global capacity): The overall traffic load does not exceed the capacity of the network.
- A2 (Local capacity): The backlog is bounded almost surely at each node.
- A3 (Contention): The waiting time in the MAC layer is bounded almost surely.

Assumption A1 is natural as we are only interested in moderately loaded networks. For a discussion on the broadcast capacity of wireless ad hoc networks see [16]. Assumption A2 means that forwarding the broadcast message together with background traffic does not locally exceed the network capacity. Assumption A3 says that once the packet is in the front of a MAC queue, it is sent in a bounded time (i.e., the neighboring nodes should not keep the channel reserved arbitrarily long).

Some comments on the interpretation of assumptions A2 and A3 are in order. These conditions would be satisfied if the local node density were bounded from above, a property that does, unfortunately, not hold in our case. Indeed, A2 and A3 are in conflict with our infinite Poisson point model for the network: Since the number of nodes in any area is a Poisson random variable and can assume arbitrarily large values, we may see arbitrarily large local node densities in a typical realization of the network. Thus, there is no bound for the number of neighbors who can claim the channel and delay a transmission.

The restrictions on the background traffic portion imposed by A2 and A3 can be interpreted as the requirement that it stays bounded independently of the underlying local node density. For example, a random point-to-point transmission per time unit in each fixed sized cell would satisfy that. As for the broadcast portion the restrictions imposed by assumptions A2 and A3 could be interpreted as a model saying that some of the broadcast packets are dropped by timeouts in the areas of

high node density. In most cases, this would not change the propagation of the broadcast at all: the neighbors of the node who dropped the broadcast packet would have neglected the packet anyway because they received it earlier from elsewhere.

E. Poisson Boolean model (PBM)

In the *Poisson Boolean model* (also known as continuum percolation), the locations of the so-called “grains” are given by the points $\{X_i\}$ of a stationary Poisson point process in \mathbb{R}^d of intensity λ (see e.g. [17], [18]). In this paper, we only consider \mathbb{R}^2 and disk shaped grains. Then we define the *occupied component* or *occupied set* $\mathcal{B}_\lambda(R)$ as the union of the grains, i.e., in our case the union of randomly scattered disks $B(X_i, R_i)$ centered at X_i and of radius R_i (see Fig. 2):

$$\mathcal{B}_\lambda(R) \doteq \bigcup_i B(X_i, R_i). \quad (3)$$

Note, that the R_i are assumed to be i.i.d., independent of the point process $\{X_i\}$, and distributed as R .




Fig. 2. Poisson Boolean model in \mathbb{R}^2 . Subcritical intensity on the left and supercritical intensity on the right.

The occupied set can be divided into disjoint clusters each of which is formed by overlapping disks. One way to measure the global connectivity is the size of the largest cluster. Let us denote $\theta_\lambda(R)$ the probability that the origin belongs to an *unbounded cluster*. Then, the *critical intensity* can be defined by

$$\lambda_c(R) \doteq \inf\{\lambda : \theta_\lambda(R) > 0\}.$$

Under quite weak¹ assumptions, the critical intensity satisfies $0 < \lambda_c < \infty$. In addition, for any *supercritical* intensity, i.e., $\lambda > \lambda_c$, the largest connected component is unbounded almost surely (a.s.). Moreover, the infinite cluster is then unique a.s. On the other hand, if the intensity is *subcritical*, i.e., $\lambda < \lambda_c$, all the connected components are finite a.s.

Neither $\theta_\lambda(R)$ nor λ_c possess known analytical expression. However, it is relatively easy to numerically estimate them. It is known, e.g., that $\lambda_c r^2 \approx 0.37$ for $R = r$ a.s. (see, e.g., [19], [20]). In particular, when setting $r = R_T/2$, we find the well-known fact that the mean number of neighbors must be larger than 4.6 in order to guarantee an unbounded communication network. For more details related to percolation in the Poisson Boolean model setting see [18].

¹For dimension $d \geq 2$, e.g., assume $\mathbb{E}[R^{2d-1}] < \infty$.

III. SPREAD OF BROADCAST FLOODING

Primary metrics reflecting broadcast performance concern how far it spreads. We discuss here two crucial metrics, one novel and one already studied: survivability and coverage. We start by defining terms and then establish useful bounds.

A. Metrics for Broadcast Spread

Since collision detection is not used in broadcast, there is no guarantee that all nodes receive the message even if the network was fully connected. Indeed, in our model carrier sensing does not protect from interference, a phenomenon called the *hidden terminal problem*, unless the carrier sensing range is sufficiently large compared to the interference range, then $R_F = R_T$ (see Fig. 1).

Collisions may result in “holes”, i.e., nodes which do not receive a broadcast packet at all, or even extinction of the broadcast. The severity of the phenomenon depends on network topology, MAC protocol and characteristics of the background traffic.

To address the severeness of collisions we study the broadcast *coverage* which denotes the long-term average proportion of the nodes which will receive the broadcast packet. We also introduce the concept of broadcast *survival* by which we mean the event that an infinite number of nodes receive the broadcast eventually. Otherwise, if only a finite number receive the message, we say that the broadcast *dies out*. Clearly, this second concept is tailored to infinite networks: in finite networks a broadcast can not survive in the above sense and, thus, the concept is not relevant. For infinite networks, however, survival constitutes a central condition without which other measures become irrelevant, as we will see.

B. Broadcast Survival

We study broadcast survival exploiting percolation theory. More precisely, we compare our channel model to a simple connection model where any nodes within a fixed, non-random *communication range* R can receive messages from each other. In particular, the model *assumes* that there are no collisions. Then, the clusters of the occupied set of the corresponding Poisson Boolean model $\mathcal{B}_\lambda(R/2)$ are exactly the set of nodes which can receive a broadcast from each other under this simple connection model. This observation leads to bounds on the survival probability of the broadcast as follows.

If the intensity λ is supercritical, i.e., $\lambda > \lambda_c(R/2)$, let $\tilde{\theta}_\lambda(R/2)$ be the probability that a random node belongs to the unbounded component (which exists and is unique almost surely). In other words, $\tilde{\theta}_\lambda(R/2)$ is the fraction of network nodes which belong to the infinite cluster. For the sake of definiteness we set $\theta_\lambda(R/2) = 0$ if $\lambda < \lambda_c(R/2)$.

Theorem 1: Assume that a broadcast is originated by a randomly picked node. Then,

$$\tilde{\theta}_\lambda(R_F/2) \leq \text{P(Broadcast Survival)} \leq \tilde{\theta}_\lambda(R_T/2). \quad (4)$$

for any background traffic and any broadcast packet length.

Proof: We start with a useful observation, as alluded above. Consider a collision-free broadcast flooding in a network defined by a Poisson Boolean model with radius $R/2$. A broadcast will survive if and only if it is originated by a node belonging to the unbounded connected component. This event has probability $\tilde{\theta}_\lambda(R/2)$.

Due to the transmission range (M2) and interference (M3), every broadcast that survives under our channel model must survive in the simple PBM connection model which *assumes* no collisions and uses the transmission range R_T as radius. In other words, a broadcast surviving under our channel model must originate in the infinite cluster of $\mathcal{B}_\lambda(R_T/2)$. The probability of this happening is the upper bound of (4).

On the other hand, due to carrier sensing (M1), each node can forward the message under our model without collisions inside a disk of radius R_F . Thus, the broadcast will survive under our channel model whenever it survives under the PBM using the collision-free radius, i.e., whenever it originates inside the unbounded component of $\mathcal{B}_\lambda(R_F/2)$, independently to traffic load and packet sizes. ■

Due to its relevance in this as well as in several later arguments we call the PBM based on the collision-free radius R_F the *collision-free PBM* and $\mathcal{B}_\lambda(R_F/2)$ the *collision-free occupied set*. Note that for this radius the transmissions from grain to grain do not experience any collision under our channel model (M1-M4) by definition of R_F .

The proof of Theorem 1 explains why the broadcast often dies out during the first steps. If a broadcast is started in a bounded component of $\mathcal{B}_\lambda(R_T/2)$, there is no path out of the component by (M2) and the broadcast can reach only small, i.e. finite number of nodes. On the other hand, consider the case where a broadcast is started in a bounded component of the collision-free occupied set $\mathcal{B}_\lambda(R_F/2)$ and assume that this component is connected to the unbounded collision-free component via nodes in $\mathcal{B}_\lambda(R_T/2)$ (the white nodes shown in Fig. 3 are such), then collisions can kill the broadcast before it reaches the collision-free unbounded component. Typically, this happens in the beginning when breaking only a few links already stops the broadcast. The longer the broadcast has survived the more likely it has reached the collision-free unbounded component of the network.

Note 1 (The case $R_F = 0$): Broadcast may survive also when $R_F = 0$ or $\lambda \in (\lambda_c(R_T/2), \lambda_c(R_F/2))$. If the broadcast packet is small, then during its retransmission there are only few “nicely” spread interfering nodes because of the carrier sensing. Thus, the number of nodes receiving the message remains large enough to help the broadcast to survive. On the other hand, if the transmitted packet is very long and there is high volume of background traffic, almost all possible sources of interference can be active during the transmission and the extinction of the broadcast is almost certain.

Note 2 (Grid networks): Assuming that the nodes are located on a grid and that there is no background traffic, it is easy to show that a broadcast will always survive under our channel model: at the “broadcast boundary” there are always receiving nodes which do not suffer from the hidden terminal

problem. Even in this simple case, holes can appear because of self-collisions of the broadcast.

C. Broadcast Coverage

Let us start by remarking that node density has a strong effect on coverage. In a dense network, there are multiple paths leading from the broadcast origin to each node. In order for a node not to receive the broadcast all of those paths have to fail due to collisions. However, one should avoid considering extremely dense networks. Indeed, 802.11 MAC cannot handle the timing if the intensity λ , and thus the average node density tends to infinity since its implementation is based on a clock of finite accuracy.

Also, coverage is closely related to survival. For meaningful results we need to assume survival of the broadcast; otherwise, if the broadcast dies out, coverage is automatically zero.

Theorem 2: Assume $\lambda > \lambda_c(R_T/2)$ and consider a random node which is picked independently of the origin of the broadcast. Then,

$$\tilde{\theta}_\lambda(R_F/2) \leq \text{P}[\text{Node covered} \mid \text{Survival}] \leq \tilde{\theta}_\lambda(R_T/2)$$

for any background traffic and any broadcast packet length.

Proof: The assumption that the broadcast does not die out means that it must have been originated inside the infinite cluster of $\mathcal{B}_\lambda(R_T/2)$. Thus the nodes located in the bounded components of $\mathcal{B}_\lambda(R_T/2)$ are certainly holes. This gives us the upper bound.

To show the lower bound, we can assume that $\lambda > \lambda_c(R_F/2)$, as otherwise $\tilde{\theta}_\lambda(R_F/2) = 0$. Then $\mathcal{B}_\lambda(R_F/2)$ has an infinite component a.s. and with probability one the broadcast will cover it: Either the broadcast is started inside the infinite component of $\mathcal{B}_\lambda(R_F/2)$ or one of the (infinitely many) links leading to the rest of the infinite component in $\mathcal{B}_\lambda(R_T/2)$ will forward the message there. Thus the potential holes are nodes which are in the bounded components of $\mathcal{B}_\lambda(R_F/2)$. ■

IV. SPEED OF BROADCAST FLOODING

Having studied broadcast performance in terms of its spread we now turn to its *dynamics*. Again we analyze two metrics, a new and an existing one, namely back-chatter and latency. As is natural in this context, we consider only the nodes which will ultimately receive the message, i.e., only the covered nodes. Moreover, our analysis focuses on asymptotical behavior so that the results hold only for a surviving broadcast.

A. Broadcast Dynamics: Metrics and Tools

We define *latency* as the time from the start of a broadcast to the first reception at a given destination node. The *duration of the back-chatter* is the time difference between the first and last reception at a given destination.

Denote by $G(\mathbf{x})$ the random set consisting of the nodes which receive the broadcast originating from the node which is nearest to the location \mathbf{x} . Notice that $G(\mathbf{x})$ depends on the underlying topology, background traffic and MAC. Let $T_F(\mathbf{x}, \mathbf{y})$ and $T_L(\mathbf{x}, \mathbf{y})$ be the first and last time, respectively,

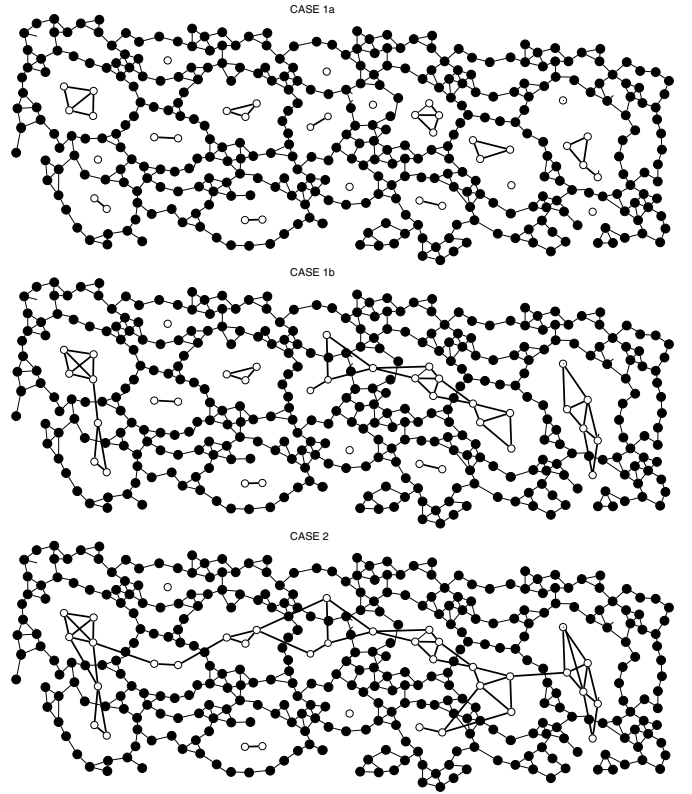


Fig. 3. The structure of the communication network. The black nodes belong to the collision-free “core” $C_\infty(\lambda, R_F)$ while the white nodes form $H(\lambda, R_T, R_F)$. The connections between white and black nodes are not shown not to overload the picture.

that the broadcast packet is received at the node in $G(\mathbf{x})$ which is nearest to the location \mathbf{y} .

Let $C(\lambda, R)$ denote the connectivity graph defined by $\mathcal{B}_\lambda(R/2)$, i.e., its vertices are the centers of the grains in $\mathcal{B}_\lambda(R/2)$ and edges are drawn between two vertices at distance at most R from each other. Let $C_\infty(\lambda, R)$ denote the unbounded component of the connectivity graph $C(\lambda, R)$.

It is useful to consider the nodes which lie in $C_\infty(\lambda, R_T)$ but not in $C_\infty(\lambda, R_F)$. Let $H(\lambda, R_T, R_F)$ be the graph formed by these nodes as vertices, with edges between two nodes that are at most at distance R_T from each other. In other words: Let $H(\lambda, R_T, R_F)$ be the subgraph of $C_\infty(\lambda, R_T)$ formed by deleting all nodes and corresponding edges that lie also in $C_\infty(\lambda, R_F)$.

We distinguish two main cases:

Case 1: $H(\lambda, R_T, R_F)$ has only bounded components;

Case 2: $H(\lambda, R_T, R_F)$ has unbounded components.

Case 1 contains two subclasses depending whether the edges of the connectivity graphs $H(\lambda, R_T, R_F)$ and $C_\infty(\lambda, R_F)$ do not cross (1a) or do cross (1b) in Fig. 3. It is an easy exercise to show that the case (1a) occurs if and only if $1 \leq R_T/R_F \leq \sqrt{3}$.

In this paper we derive results mainly for Case 1. Case 2 is left for future’s research. We conjecture that there is a phase transition between Case 1 and Case 2, i.e., there exists a critical

intensity $\lambda_H \doteq \lambda_H(R_T, R_H)$ and a corresponding critical transmission range $R_H \doteq R_H(\lambda, R_F)$ such that whenever $\lambda > \lambda_H$ (or $R_T > R_H$) Case 2 occurs almost surely. Exploring such issues is beyond the scope of this paper.

Intuitively speaking, we may think of the nodes in $C_\infty(\lambda, R_F)$ as forming the main collision-free communication core, where communication is guaranteed not to suffer from collision due to carrier sensing (M1). The nodes in this core are connected over a distance less than R_F , but under suitable traffic condition the communication distance can be up R_T . The nodes and components of $H(\lambda, R_T, R_F)$ can then be thought of as some kind of sideroads, sometimes short-cutting and speeding the broadcast up, but always without any service guarantees.

B. Latency

Let us first consider latency in a collision-free network.

Lemma 1: Assume A1–A3, a fixed transmission range R and an idealistic MAC without collisions. If $\lambda > \lambda_c(R/2)$ and the nearest node to the origin belongs to $C_\infty(\lambda, R)$, then for a fixed \mathbf{y} with $|\mathbf{y}| = 1$ there exists finite strictly positive constants $\eta(\lambda, R)$ and $\gamma(\lambda, R)$ such that

$$\gamma(\lambda, R) \geq \frac{T_F(\mathbf{0}, a\mathbf{y})}{|a\mathbf{y}|} \geq \eta(\lambda, R) \quad a.s.$$

whenever a is large enough.

Note that the constants $\gamma(\lambda, R)$ and $\eta(\lambda, R)$ are different if the backlog and contention bounds are changed.

Proof: This is a straightforward modification of a proof in [21, Thm. 1]. To get a lower bound, we assume that neither backlog nor contention exist. Then the packet is forwarded instantaneously and the delay at each node is the (constant) transmission time. For an upper bound, we can assume that the packet stays at each node the maximal backlog and contention durations, which both are bounded almost surely. ■

The next theorem is the main result of this section. It gives conditions under which the propagation speed of flooding is asymptotically linear.

Theorem 3: Assume A1–A3 and Case 1. If a broadcast originating from the nearest point to the origin does not die out and $\lambda > \lambda_c(R_F/2)$, then there exist finite strictly positive constants $\eta = \eta(\lambda, R_T, R_I, R_C)$ and $\gamma = \gamma(\lambda, R_T, R_I, R_C)$ such that for a fixed \mathbf{y} with $|\mathbf{y}| = 1$

$$\gamma \geq \frac{T_F(\mathbf{0}, a\mathbf{y})}{|a\mathbf{y}|} \geq \eta \quad a.s., \quad (5)$$

whenever a is large enough.

Proof: (Sketch) Upper bound. Consider the Poisson Boolean model $\mathcal{B}_\lambda(R_F/2)$ and assume that the source (the nearest node to the origin) and the destination (the nearest node to $a\mathbf{y}$ in $G(\mathbf{0})$) are not in $C_\infty(\lambda, R_F)$ (the collision free unbounded component). Then the random variable $T_F(\mathbf{0}, a\mathbf{y})$ can be bounded by a sum of the following three terms (see Fig. 4):

- 1) *Time until the nearest node to the origin in the unbounded collision free component receives the message.*




Fig. 4. Constructing a bound for latency: the black nodes form the unbounded cluster of $\mathcal{B}_\lambda(R_F/2)$ and white nodes belong to the unbounded cluster of $\mathcal{B}_\lambda(R_T/2)$. The connections between black and white nodes are not shown in the picture.

In Case 1, the number of nodes of $H(\lambda, R_T, R_F)$ which can forward the message until it enters $C_\infty(\lambda, R_F)$ is finite. The Euclidean distance between the point where the message enters $C_\infty(\lambda, R_F)$ and location of the nearest node to the origin in $C_\infty(\lambda, R_F)$ is thus finite. By finite delays at each node and [21, Prop. 4]), it takes a finite time until the nearest node of $C_\infty(\lambda, R_F)$ to the origin receives the message.

- 2) *Time to “travel” from the neighborhood of the origin to the neighborhood of $a\mathbf{y}$ on $C_\infty(\lambda, R_F)$.* Lemma 1 describes asymptotics of the latency between the nodes in the unbounded collision free component nearest to the origin and the node nearest to point $a\mathbf{y}$.
- 3) *Time until the message propagates from the collision free component to the destination.* Analogously to 1, all the exit points from $C_\infty(\lambda, R_F)$ to the bounded component of $H(\lambda, R_T, R_F)$ containing the destination node are located at finite distances from the nearest point to $a\mathbf{y}$ in $C_\infty(\lambda, R_F)$.

All the above reasonings are with probability one. As duration of 1) and 3) are finite (and asymptotically independent), only 2) matters asymptotically.

Lower bound. By the assumption of a surviving broadcast, the nearest node to the origin belongs to $C_\infty(\lambda, R_T)$. Because of collisions, it may happen that the destination node is not the nearest node to $a\mathbf{y}$ in $C_\infty(\lambda, R_T)$. Analogous reasoning as in above, based on boundedness of the components in $H(\lambda, R_T, R_F)$, shows that asymptotically this does not matter and a lower bound can be found by considering an ideal MAC over $C_\infty(\lambda, R_T)$ as in the proof of Lemma 1. ■

The theorem shows that in the worst case the broadcast can spread only by small steps which are guaranteed by the collision free radius R_F . This can happen even under light load if the background traffic is bad enough. If the delays due to contention and background traffic in each node are approximately (or even asymptotically) independent, a reasoning invoking the generalized law of large numbers shows that increasing the traffic load makes both bounds larger. This also seen in the simulations.

Note 3: If we change setting in such way that source and destination pairs are always in $C_\infty(\lambda, R_F)$, then Theorem 3 holds true also in Case 2.

Note 4: Latency has an ‘‘asymptotic independence of the density’’ if no background traffic is present. Due to carrier sensing, the density of active nodes at a time point cannot be arbitrary high. Even in medium dense networks, most of the nodes are idle.

C. Back-chatter

Back-chatter period is the time interval between the first and last hear of the same broadcast message at a given location. The length of the back-chatter is defined

$$T_B(\mathbf{x}, \mathbf{y}) \doteq T_L(\mathbf{x}, \mathbf{y}) - T_F(\mathbf{x}, \mathbf{y}).$$

The number of times a broadcast message received by a node is bounded by the number of neighbors. Due to background traffic and collisions it could happen that some of the neighboring nodes receive the message much later. A direct application of Theorem 3 implies that the duration of back-chatter cannot increase faster than linearly, i.e., under Case 1

$$\lim_{a \rightarrow \infty} \frac{T_B(\mathbf{0}, a\mathbf{y})}{|a\mathbf{y}|} \leq \gamma - \eta \quad \text{a.s..}$$

The following theorem shows that the linear growth bound is overly pessimistic and the length of back-chatter is bounded by a finite random variable.

Theorem 4: Assume A1–A3, Case 1, $\lambda > \lambda_c(R_F/2)$ and a surviving broadcast flooding. Then there exist a finite random variable $M \doteq M(\lambda, R_T, R_F)$ such that for any \mathbf{x} and \mathbf{y} in \mathbb{R}^2 ,

$$T_B(\mathbf{x}, \mathbf{y}) \stackrel{\text{distr}}{\leq} M.$$

Proof: Consider first an arbitrary node Y and the (bounded) components of $H(\lambda, R_T, R_F)$ that are at most R_T away from node Y . Let B_Y denote the union of the nodes in these components and $\#B_Y$ the number of nodes in B_Y . The surrounding $C_\infty(\lambda, R_F)$ nodes are contained in set

$$A_Y = \{z \in C_\infty(\lambda, R_F) : \exists w \in B_Y \cup Y \text{ s.t. } |z - w| \leq R_T\}$$

If all nodes of A_Y have retransmitted the broadcast packet, we can be sure that the message cannot enter later in the bounded components of $H(\lambda, R_T, R_F)$ in the neighborhood of Y (which is B_Y).

If a node Y in $C_\infty(\lambda, R_F)$ has received a broadcast message, then any finite subset of $C_\infty(\lambda, R_F)$ will hear the message in a finite time almost surely. Assuming that the broadcast starts at Y , communication occurs only on $C_\infty(\lambda, R_F)$ and each transmission takes the maximum time (d_{\max}), then we denote the time till all nodes in A have received and retransmitted the message by $S(Y, A)$. Note that some paths may need to visit outside of set A .

If node Y in $H(\lambda, R_T, R_F)$ has received the broadcast message, then one of the surrounding $C_\infty(\lambda, R_F)$ nodes must have received it earlier, because the destination node is in an




Fig. 5. Survivability in a sparse network with high background load. Simulations are done with 700 nodes in an area of 4 km \times 4 km.

isolated component of $H(\lambda, R_T, R_F)$. We denote this node by \tilde{Y} . Clearly $B_{\tilde{Y}} \supset B_Y$ and $A_{\tilde{Y}} \supset A_Y$.

Let Y_y be the nearest node to \mathbf{y} . If Y_y in $C_\infty(\lambda, R_F)$, we denote $Y_y = \tilde{Y}_y$. Then

$$\begin{aligned} T_B(\mathbf{x}, \mathbf{y}) &\leq S(Y, A_{\tilde{Y}_y}) + d_{\max} \#B_{\tilde{Y}_y} \\ &\stackrel{\text{distr}}{=} S(0, A_{\tilde{Y}_0}) + d_{\max} \#B_{\tilde{Y}_0} = M, \end{aligned}$$

where M is a finite random variable. \blacksquare

V. SIMULATION STUDIES

We employ the NS-2.29 simulator. Nodes are uniformly distributed in a 4 km by 4 km area and each node has the default radio range of 250 m. We use the IEEE 802.11b MAC protocol with basic transmission rate 1Mbps and broadcast packet size of 125 Bytes. The originating node of broadcast flooding is in the center of square area. The background traffic is generated by ‘‘Hello’’ messages with size of 125 Bytes.

In the simulations, we consider three different node densities (700, 1600 and 2700 nodes in the square) corresponding to the mean number of neighbors 8.6, 19.6 and 33.1. Moreover, we have implemented two background traffic scenarios, light and high. When changing the density of the network, we keep the volume of the background traffic per area fixed. Thus the capacity left for the broadcast does not change. However, in the denser networks, flooding itself consumes more capacity.

Unfortunately, we were not able to dig into the details of the ns-2 code and find out how to estimate the interference and carrier sensing range. Thus the following simulations demonstrate only the qualitative behavior implicated by the analytical results.

In Fig. 5, surviving broadcasts are denoted by the bars in the right and those which died out correspond to the bars near the origin. In the simulation experiment, broadcast survived approximately in 90% of cases. This could be compared with the fraction of nodes in the infinite cluster in the Poisson




Fig. 6. Latency as a function of Euclidean distance. Simulations with light and high background loads are shown in each picture.

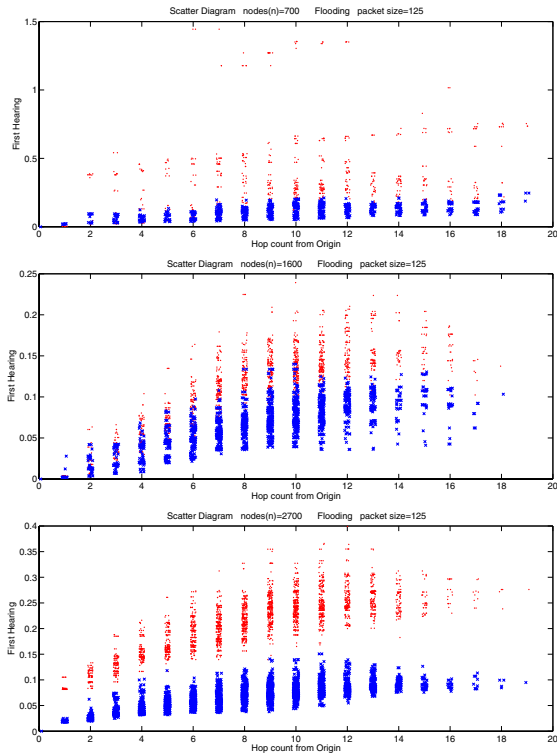


Fig. 7. Latency as a function of hop count. Simulations with light and high background loads are shown in each picture.




Fig. 8. Duration of back-chatter as a function of distance.

Boolean models $\mathcal{B}_\lambda(R_T/2)$ and $\mathcal{B}_\lambda(R_F/2)$, with $\lambda = 43.75$ and $R_T = 0.25$. By simple numerical simulations one finds that $\tilde{\theta}_\lambda(R_T/2) > 0.99$. If the background is very low the survivability probabilities are very close to one (not shown here). Whereas with high background traffic (shown in Fig. 5) the approximation $\tilde{\theta}_\lambda(R_F/2)$ should be used. As mentioned earlier, we do not know the guaranteed communication range R_F in the ns-2 simulator. By reverse engineering, we estimate that $R_I \leq 0.8R_T$ in this case.

Fig. 6 shows how latency depends on density and background traffic. The asymptotic linear bounds predicted by Theorem 3 are clearly seen in the scatter plots. This becomes especially clear in the two lower pictures in 6 where the density is higher. The qualitative behavior remains the same even if latency is plotted as a function of hop count as seen in Fig. 7. This is no surprise because Euclidean distance predicts quite accurately the number of hops of a shortest path between two nodes.

Back-chatter is considered only in the two densest networks. The simulations illustrated in Fig. 8 and 9 demonstrate that the length of the back-chatter period stays bounded. However, one should notice that due to boundary effects the difference between the first and last reception times could be poorly estimated because in the simulation setting there is no possibility that the broadcast “returns” from outside the $4 \text{ km} \times 4 \text{ km}$ square.

VI. CONCLUDING REMARKS

We have derived results which enable an informed decision whether broadcast via flooding is reasonable solution in a wireless ad hoc network. Assuming that the capacity of the network is neither globally nor locally exceeded by retransmissions and background traffic, flooding usually matches (or even outperforms) other broadcast mechanisms with respect to survivability and coverage. Moreover, the asymptotic latency for flooding is only a constant order from the optimal: indeed it




Fig. 9. Duration of back-chatter as a function of hop count.

is well-known that a message cannot spread faster than linearly in geometric networks.

Back-chatter constitutes an obvious waste of bandwidth. The total back-chatter load, i.e., number of needless retransmissions, naturally remains the same everywhere due to the assumed homogeneity of the network. In principle, the farther from the origin of a broadcast, the longer the back-chatter can last. Under some assumptions on the system parameters, our analysis shows that the length of a back-chatter period stays stochastically bounded independently of the distance to the origin.

In addition to characterizing broadcast flooding, the results of this paper are useful for other broadcast mechanisms. If a backbone is used for broadcast, e.g., then already a few collisions may have a severe impact. Since such backbones are optimized with respect to redundancy, the retransmitting nodes form a very sparse network which is — in the absence of collisions detection — vulnerable to collisions caused by other traffic. Sacrificing some redundancy could increase robustness against collisions. A rule of thumb could be to choose a backbone of nodes such that message can spread along connections of length at most R_F .

In this paper, in addition to describing the qualitative behavior, we have presented the best and worst case performance bounds for broadcast flooding. A natural extension of the current work would be to show how the averaging properties based on ergodicity or independence assumptions on the delays at nodes could be used to improve the bounds. Unfortunately, one of the problems with percolation based analysis is that most of the results are only about existence and the analytical formulae are rarely known.

ACKNOWLEDGMENT

Work was done when P. Mannersalo was visiting Rice University. Supported in part from NSF, grant number ANI-0338856, and Texas ATP, project number 003604-0036-2003.

REFERENCES

- [1] X. Hong, K. Xu, and M. Gerla, "Scalable routing protocols for mobile ad hoc networks," *IEEE Network*, vol. 16, pp. 11–21, 2002.
- [2] S.-Y. Ni, Y.-C. Tseng, Y.-C. Chen, and J.-P. Sheu, "The broadcast storm problem in a mobile ad hoc network," in *Proceedings of MobiCom 1999*. ACM Press, 1999, pp. 151–162.
- [3] B. Williams and T. Camp, "Comparison of broadcasting techniques for mobile ad hoc networks," in *Proceedings of MobiHoc 2002*. ACM Press, 2002, pp. 194–205.
- [4] W. Lou and J. Wu, "Localized broadcasting in mobile ad hoc networks using neighbor designation," in *Mobile Computing Handbook*, M. Ilyas and I. Mahgoub, Eds. CRC Press, 2005, ch. 28.
- [5] J. Wu and F. Dai, "Performance analysis of broadcast protocols in ad hoc networks based on self-pruning," *IEEE Trans. Parallel Distrib. Syst.*, vol. 15, no. 11, pp. 1027–1040, 2004.
- [6] IEEE-SA Standards Board, "IEEE Standard for Information Technology: Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE Computer Society LAN MAN standards Committee, 1999.
- [7] S. Alagar, S. Venkatesan, and J. Cleveland, "Reliable broadcast in mobile wireless networks," in *Proceedings of IEEE Military Communications Conference 1995*, 1995, pp. 236–240.
- [8] K. Tang and M. Gerla, "MAC reliable broadcast in ad hoc networks," in *Proceedings of IEEE Military Communications Conference 2001*, vol. 2, 2001, pp. 1008–1013.
- [9] M.-T. Sun, L. Huang, A. Arora, and T.-H. Lai, "Reliable MAC layer multicast in IEEE 802.11 wireless networks," in *Proceedings of the 2002 International Conference on Parallel Processing (ICPP'02)*. IEEE Computer Society, 2002.
- [10] J. Chen and M. Huang, "A broadcast engagement ack mechanism for reliable broadcast transmission in mobile ad hoc networks," *IEICE Transactions on Communications*, vol. E88-B, no. 9, pp. 3570–3578, 2005.
- [11] K. Viswanath and K. Obraczka, "Modeling the performance of flooding in multihop ad hoc networks (extended version)," *Computer Communications Journal (CCJ)*, 2005.
- [12] O. Dousse, P. Thiran, and M. Hassler, "Connectivity in ad hoc and hybrid networks," in *Proceedings of IEEE Infocom 2002*. IEEE Computer Society, 2002.
- [13] L. Booth, J. Bruck, M. Franceschetti, and R. Meester, "Covering algorithm, continuum percolation and the geometry of wireless networks," *Ann. Appl. Probab.*, vol. 13, no. 2, pp. 722–741, 2003.
- [14] O. Dousse, M. Franceschetti, and P. Thiran, "Information theoretic bounds on the throughput scaling of wireless relay networks," in *Proceedings of IEEE Infocom 2005*. IEEE Computer Society, 2005.
- [15] O. Dousse, F. Baccelli, and P. Thiran, "Impact of interferences on connectivity in ad hoc networks," in *Proceedings of IEEE Infocom 2003*. IEEE Computer Society, 2003.
- [16] A. Keshavarz-Haddad, V. Ribeiro, and R. Riedi, "Broadcast capacity in multihop wireless networks," in *Proceedings of MobiCom 2006*. ACM Press, 2006.
- [17] D. Stoyan, W. Kendall, and J. Mecke, *Stochastic geometry and its applications*, 2nd ed. Chichester: Wiley, 1995.
- [18] R. Meester and R. Roy, *Continuum percolation*, ser. Cambridge tracts in mathematics. Cambridge University Press, 1996, vol. 119.
- [19] T. Vicsek and J. Kertsz, "Monte Carlo renormalisation group approach to percolation on a continuum: Test of universality," *J. Phys. A*, vol. 14, no. L31, 1981.
- [20] J. Quintanilla, S. Torquato, and R. Ziff, "Efficient measurement of the percolation threshold for fully penetrable discs," *J. Phys. A*, vol. 33, pp. 399–407, 2000.
- [21] O. Dousse, P. Mannersalo, and P. Thiran, "Latency of wireless sensor networks with uncoordinated power saving mechanisms," in *Proceedings of Mobihoc*, 2004, pp. 109–120.