

Controlled Gossiping in Ad Hoc Networks

Ivan S. Lysiuk
Genesys Telecommunication Laboratories
Alcatel-Lucent
2001 Junipero Serra Blvd
Daly City, CA 94014, USA

Zygmunt J. Haas
Wireless Networks Lab
School of Electrical and Computer Engineering
Cornell University
Ithaca, NY 14853, USA

Abstract—Gossiping, a probabilistic algorithm for network-wide broadcasting, can significantly reduce the number of transmissions in one-to-all message spread. Based on the *Continuum Percolation Theory*, we develop a novel analytical model of gossiping for wireless *ad hoc* networks. This model provides an intuitive way for setting the gossiping probability, so that a gossip achieves high degree of network coverage with limited number of transmissions. We also propose and study a distributed gossiping scheme, in which nodes set their gossiping probabilities according to locally acquired information. Both gossiping schemes are demonstrated to be scalable; i.e., for fixed network area, the expected number of gossip transmissions does not depend on the number of nodes in the network. The analytical results are verified via simulations of large *ad hoc* networks.

I. INTRODUCTION

Mobile *ad hoc* networks are networks with no pre-existing infrastructure, where nodes communicate over the wireless medium using multi-hop routing schemes.

A network-wide broadcast operation, in which a single node sends a message to all other nodes of the network, is a fundamental operation in *ad hoc* networks and is used, for example, for path discovery by such routing protocols as AODV, DSR, and ZRP [1]. The simplest way to perform network-wide broadcast is through flooding. Assuming no transmission errors, flooding guarantees delivery of the message to all nodes connected to the message originating node. However, flooding is also quite inefficient in the number of transmissions of the message. Furthermore, due to the excessive transmissions, flooding leads to increased contention in the network – a problem known as the *broadcast storm* [2].

The probabilistic technique of broadcasting in MANETs, termed *gossiping*,¹ was initially suggested in the work by Ni et al [2], where each node retransmits the broadcast message to all its neighbors with some fixed probability. A similar approach was studied in the context of its application to routing in *ad hoc* networks by Haas et al [3]. In [4], the authors used the term *probabilistic broadcast* and, similarly to [3], pointed out that the gossiping performance has very strong signs of percolation behavior.

Despite its natural simplicity, gossiping as a network-wide broadcasting algorithm was not given sufficient consideration in the technical literature. One reason is the absence of rigorous analytical models which are able to relate the gossiping

probability to the performance of gossiping schemes, as a function of network parameters. This paper fills in this gap and significantly extends the analytical results of [3], where *ad hoc* networks were modeled as *infinite regular graphs*. Instead, we more realistically model *ad hoc* networks as *Random Plane Graphs* [5]. Our work makes the following contributions.

1) We develop a new percolation model of gossiping, which answers a key problem: for an *ad hoc* network with given average node degree ϕ and for any coverage level $0 < \gamma < 1$, we show how to set the gossiping probability p_g , so that the probability that the gossip spreads over the entire network is γ . Furthermore, in each such network-wide gossip spread, each node receives the gossip with probability γ , while the expected number of gossip transmissions is as small as possible.

2) We design and analyze another scheme, the *Distributed Gossip* algorithm, with performance similar to that of the *Global Gossip*, but without requiring the knowledge of the average node degree. In *Distributed Gossip*, each node is required to know only the number of its direct neighbors. We find sufficient conditions for *Distributed Gossip* to achieve any coverage level $0 < \gamma < 1$.

3) We show that both of the gossiping schemes are scalable in the following sense: given a large network with sufficiently large average node degree, the expected number of transmissions in each gossip spread is equal to (in the *Distributed Gossip*, it is upper-bounded by) a value, which behaves as a sub-linear function of the network area and which, in particular, does not depend on the number of nodes in the network.

The organization of the paper is as follows: Section II summarizes some pertinent related work. Section III derives relevant results from the *Continuum Percolation Theory* and studies their applicability to the case of finite-size networks. Section IV presents the *Global Gossip* algorithm and its analytical model. The *Distributed Gossip* scheme is introduced and evaluated in Section V, while in Section VI, we experimentally evaluate the two gossiping schemes. Finally, Section VII concludes the work.

II. RELATED WORK

A. Network-Wide Broadcasting and Gossiping

The paper by Williams and Camp [6] provides a comprehensive comparative analysis of the existing broadcasting protocols for *ad hoc* networks: twelve broadcasting schemes

¹In this paper, we also refer to this scheme as *Global Gossip*, to be distinguished from the *Distributed Gossip*, which we introduce in Section V.

are divided into four categories, and the performance of the selected protocols in each category is studied.

Another broad review of different broadcast approaches is presented in [7] by Stojmenović and Wu. The paper provides a taxonomy of broadcasting methods based on their determinism, type of network information required, reliability, and contents of the “hello” and broadcast messages.

The works by Haas et al [3] and by Sasson et al [4] explore the percolation properties of gossiping and are directly related to our work in this paper. The authors of [3] were the first to link gossiping to the percolation theory and to note the bimodal behavior of gossiping. Also, a number of heuristic enhancements of gossiping were proposed in [3], and gossiping was shown to yield significant savings of transmissions, when applied to AODV routing.

Cartigny and Simplot [8] suggested and empirically studied four probabilistic broadcasting schemes, in which nodes utilize the knowledge of their local density and neighboring nodes can exchange their neighborhood lists between themselves. One of the schemes studied (mode 2) is similar to our *Distributed Gossip* scheme. However, [8] relies on simulation for comparison of the performance of the proposed schemes, while our work is primarily based on analytical results.

B. Continuum Percolation

Since the seminal paper by Gilbert [9], which introduced the topic of *Continuum Percolation*, random geometric graphs have been extensively studied and widely used as a model for radio networks and, lately, for *ad hoc* networks as well. In this section, we review some percolation properties of random plane graphs and discuss the properties which we later use for the derivation of our analytical gossiping model for *ad hoc* networks.

A random geometric graph is defined by placement of points of a homogeneous Poisson process \mathcal{H}_λ of intensity λ in \mathbb{R}^2 , where each two points of the process \mathcal{H}_λ are connected by an edge if the distance between them is not greater than r . The set of points \mathcal{H}_λ and the set of edges between them form a *random plane graph* $G(\mathcal{H}_\lambda, r)$. As in [5], we consider a process $\mathcal{H}_{\lambda,0} \triangleq \mathcal{H}_\lambda \cup \{\mathbf{0}\}$, where $\mathbf{0}$ is the origin of \mathbb{R}^2 . Considering the associated graph $G(\mathcal{H}_{\lambda,0}, 1)$, the percolation probability $p_\infty(\lambda)$ is the probability that the node at the origin $\mathbf{0}$ is connected to infinitely many nodes of the graph $G(\mathcal{H}_{\lambda,0}, 1)$. The percolation threshold density is $\lambda_c = \inf\{\lambda > 0 : p_\infty(\lambda) > 0\}$. The percolation behavior is as follows: in the sub-critical regime (i.e., $\lambda < \lambda_c$), the nodes in $G(\mathcal{H}_{\lambda,0}, 1)$ form numerous components² of finite size, whereas in the super-critical regime (i.e., $\lambda > \lambda_c$), there exists at least one infinitely large component in $G(\mathcal{H}_{\lambda,0}, 1)$ and the origin belongs to an infinitely large component with probability $p_\infty(\lambda) > 0$.

As discussed in [10], [11], and [9], the critical value λ_c is a key parameter when connectivity is considered. The bounds on λ_c were analytically estimated in [12], and the value of λ_c

was experimentally evaluated by Torquato [13]. In our work, we are interested in the behavior of the probability $p_\infty(\lambda)$ when $\lambda > \lambda_c$.

III. PRELIMINARIES: PERCOLATION IN FINITE GRAPHS

This section lays the basis toward obtaining our models of gossiping. First, considering an *ad hoc* network as a random plane graph, we argue that there exists an average node degree that corresponds to a given fraction of nodes in an infinite connected component³. Next, we study finite-size networks with random uniform node distribution, and we experimentally investigate the impact of finiteness on their percolation properties. The results obtained in this section are used in our analytical models, as well as in simulations of the gossiping schemes, studied in the following sections.

A. Continuum Percolation in Infinite Graphs

The transmission radii of all nodes in our *ad hoc* network are assumed to be the same, and all links are assumed to be bidirectional; i.e., any two nodes located at a distance not greater than r are able to directly communicate with each other. Consequently, the network model considered in our formal analysis is a random plane graph $G(\mathcal{H}_\lambda, r)$. Furthermore, the scaling property of random geometric graphs ([5], [12]) allows us to merge the two parameters λ and r of our random plane graph into a single parameter $\phi = \lambda\pi r^2$. Thus, we denote a random graph $G(\mathcal{H}_\lambda, r)$ as $G(\phi)$, where $\phi = \lambda\pi r^2$, and we consider the percolation probability $\theta(\phi) = p_\infty(\lambda)$ as a function of ϕ , with the critical average node degree defined as $\phi_c = \lambda_c\pi r^2$.

According to Penrose [5, Theorem 9.19], in the super-critical regime, almost surely there is only a single unbounded component in $G(\phi)$. Hence, for some realization of $G(\phi)$ for which we assume that an infinite component exists in $G(\phi)$, we refer to it as the unbounded component. Thus, the probability $\theta(\phi)$ could be interpreted in the case of $\phi > \phi_c$ as the probability that an arbitrary node of \mathcal{H}_λ belongs to the unbounded component of $G(\phi)$. When $\phi < \phi_c$, no unbounded component exists in $G(\phi)$, and hence $\theta(\phi) = 0$.

We notice that for $\phi > \phi_c$, the function $\theta(\phi)$ is continuous and non-decreasing. Furthermore, $\theta(\phi) \rightarrow 1$ as $\phi \rightarrow \infty$ [5, Theorem 9.20 and Proposition 9.21]. Therefore we define its inverse function $\theta^{-1}(\beta)$ as follows: for $0 < \beta < 1$, let $\theta^{-1}(\beta) = \min\{\phi : \theta(\phi) = \beta\}$.

Having defined $\theta^{-1}(\beta)$, we have to remark that there is no known analytical expression neither for $\theta(\phi)$, nor for $\theta^{-1}(\beta)$. Therefore, in our analysis of the gossiping algorithm, we just assume the existence of $\theta^{-1}(\beta)$. Nevertheless, for the practical cases of finite-size networks, we estimate the behavior of $\theta^{-1}(\beta)$ in subsection III-B.

B. Finite graphs

In this work, we consider the practical case of large but finite networks that could undergo a percolation transition. In our model, N nodes of an *ad hoc* network are uniformly and

²By component we mean a connected sub-graph.

³We also term it an *unbounded component*.




Fig. 1. Experimental evaluation of $\theta_\alpha(\phi)$ for various network sizes

randomly distributed over a region $[-L/2, L/2] \times [-L/2, L/2]$ with the area of $A = L^2$. For $A \gg \pi r^2$, this placement of nodes could be well approximated as a homogeneous Poisson process \mathcal{H}_λ with intensity $\lambda = \frac{N}{A}$, restricted to the region $[-L/2, L/2] \times [-L/2, L/2]$. For this case, it was shown by Penrose [5] that the fraction of nodes in the largest cluster converges in probability to $\theta(\lambda)$ as $L \rightarrow \infty$, where the node density is λ . This result is supported by the experiments in [9] and [11], which demonstrated that the fraction of nodes in the largest connected component is a non-decreasing function, rapidly growing in the proximity of the critical density λ_c .

We experimentally estimate the fraction of nodes in the largest connected component, which we denote as $\theta_\alpha(\phi)$, for networks of different sizes, where $\alpha = \frac{L}{r}$ defines the network size in units of the transmission radius. The experimental evaluation of the behavior of $\theta_\alpha(\phi)$ with $\phi = \pi r^2 \frac{N}{A}$ is presented in Fig. 1. We obtained these results by varying the number of nodes in fixed-area networks and measuring the fraction of nodes in the largest connected cluster.⁴ As expected, the percolation transition happens in finite networks, with the steepest part occurring in the proximity of $\phi_c \approx 4.52$ (also see [13]). When the average node degree is above this threshold, the percolation in the networks of smaller size is “smoother” due to the fact that a greater fraction of nodes lays close to the boundary of the network.

We are interested in the behavior of $\theta_\alpha^{-1}(\beta)$ as a function of α for fixed values of β in the proximity of 1, which is shown in Fig. 2. Recall that $\theta_\alpha^{-1}(\beta)$ is the minimal average degree ϕ such that the fraction of nodes in the largest connected component of a network with $L = r\alpha$ is at least β . The values of β considered in Fig. 2 were selected based on the following practical considerations. On one hand, *Continuum Percolation Theory* does not guarantee that the percolation probability reaches 1 for any finite value of average node degree. On the other hand, as the conducted experiments demonstrate, practical values of β (slightly less than 1) are achievable for finite average node degree.

⁴The vertical bars in the figures denote 95% confidence intervals.

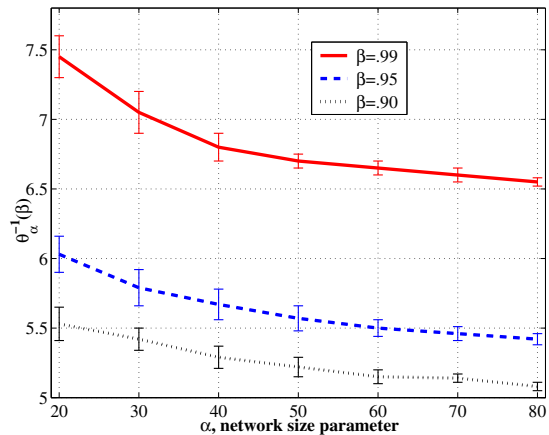


Fig. 2. Experimental evaluation of $\theta_\alpha^{-1}(\beta)$ as a function of network size α for different values of β

IV. GLOBAL GOSSIP

A. Description

We assume a targeted gossip *coverage level*, $0 < \gamma < 1$, and the average degree of network nodes, $\phi \geq \theta^{-1}(\gamma)$. The following gossiping algorithm is analyzed in this section. The gossip source node broadcasts the gossip message to all of its neighbors. Any node, except the gossip originator, upon receiving the gossip message for the first time rebroadcasts it with probability

$$p_g = \theta^{-1}(\gamma)/\phi. \quad (1)$$

When a node receives the gossip message after it had already received the message before, the node ignores the message. We refer to this scheme as *Global Gossip*.

B. Analytical Model

To analyze the scheme described in Section IV-A, the following equivalent scenario of a single gossip spread is considered. First, each node independently, with probability p_g , decides *a priori* whether it is going to participate in the gossip spread. Thus, nodes randomly choose to be either *active* or *passive* gossipers. Next, an initiator of the gossip is randomly chosen among all the nodes. Finally, the gossip initiator, regardless of its *a priori* decision on gossip participation, broadcasts the gossip message to all its neighbors, and each *active* node does so if and when it receives the gossip.

The equivalent scenario described above is inspired by the *principle of deferred decisions* in [14]. Our approach can be termed a *principle of a priori decisions*, which is the only non-intuitive step in our analysis of gossiping. This method lets us study gossiping, which is a dynamic process, in a time-independent manner by analyzing percolation properties of a random plane graph formed on the set of *active* nodes S_a .

We define a *successful* gossip spread in the following manner: a gossip run is *successful* if the gossip message was received by infinitely many nodes. However, in an infinite network there will also be infinitely many nodes that do not receive the gossip. Therefore, the quantity of interest in a

successful gossip spread is the fraction of nodes receiving the gossip.

Theorem 1: When nodes are gossiping with probability $p_g = \frac{\theta^{-1}(\gamma)}{\phi}$ in a network with average node degree $\phi \geq \theta^{-1}(\gamma)$, then:

- (a) each gossip run is successful with probability γ ;
- (b) in the case of a successful gossip run, an arbitrary node receives the gossip message with probability γ .

Proof: Since the node placement is Poisson with intensity λ , and each node belongs to S_a with probability p_g independently, then the nodes in S_a are distributed according to a Poisson process with intensity λp_g . This implies the average node degree of $\lambda p_g \pi r^2 = \theta^{-1}(\gamma)$ in the random plane graph of nodes in S_a . Therefore, an active node belongs to the unbounded component of active nodes with probability γ .

Let's denote the set of active nodes in the unbounded component of active nodes as S_{a+} and the set of active nodes that do not belong to the unbounded component as S_{a-} . Similarly, we define (disjoint) subsets of S_p ; S_{p+} is the set of passive layer nodes with at least one neighbor from S_{a+} , and $S_{p-} = S_p \setminus S_{p+}$.

We find probabilities that a node belongs to the different subsets of active and passive layers and show that any node (passive or active) either belongs to the unbounded component of active nodes or is adjacent to it with probability γ .

The probability that a node n belongs to the set S_{a+} is:

$$P(n \in S_{a+}) = P(n \in S_a)P(n \in S_{a+}|n \in S_a) = p_g \gamma. \quad (2)$$

Due to spacial invariance of the Poisson process on S_a , the origin of the coordinates $\mathbf{0}$ can be placed anywhere, including at any passive node. Thus, the probability that a passive node is connected to the unbounded cluster of active nodes S_{a+} is:

$$P(n \in S_{p+}|n \in S_p) = \gamma.$$

Correspondingly,

$$P(n \in S_{p+}) = (1 - p_g)\gamma. \quad (3)$$

We define the set $S_+ = S_{a+} \cup S_{p+}$. A node in S_+ either belongs to the unbounded component of the active layer, or the node has at least one neighbor that belongs to the unbounded component of the active layer. As follows from equations (2) and (3):

$$P(n \in S_+) = P(n \in S_{a+}) + P(n \in S_{p+}) = \gamma.$$

If a gossip initiator belongs to S_+ , the gossip will be received by infinitely many nodes of the set S_+ . Since a gossip originator belongs to S_+ with probability γ , a gossip run is *successful* with probability γ . If the gossip is *successful*, all nodes in S_+ receive it. Consequently, an arbitrary node n will receive the gossip with probability $P(n \in S_+) = \gamma$. ■

The probability that an arbitrary node transmits a gossip in a successful run is $\gamma \frac{\theta^{-1}(\gamma)}{\phi}$, which is a strictly increasing function of γ for fixed value of ϕ . Therefore, it is impossible to lower the fraction of transmitting nodes without lowering the expected fraction of nodes receiving the gossip, when all the nodes gossip with the same gossip probability.

C. Scalability of the Global Gossip

When applying the above results to finite networks of N nodes distributed uniformly over a square area A , $\theta_\alpha^{-1}(\gamma)$ should be used instead of $\theta^{-1}(\gamma)$, so that the finiteness of the network is accounted for. In this case, the expected number of gossip transmissions per single successful gossip spread is:

$$E[N_{tx}] = N\gamma p_g = N\gamma \frac{\theta_\alpha^{-1}(\gamma)}{\phi} = \gamma \theta_{\frac{\sqrt{A}}{r}}^{-1}(\gamma) \frac{A}{\pi r^2}. \quad (4)$$

As follows from equation (4), the Global Gossip algorithm is scalable in the following sense: for networks of a fixed area, the expected number of transmissions in the case of a successful gossip does not depend on the number of network nodes, but rather on the network area.

V. DISTRIBUTED GOSSIP

A. Description

The *Global Gossip* scheme that we have presented in Section IV is based on the assumption that the average node degree is known to all nodes and that its value does not change with time. These can be quite restricting assumptions and could be considered unacceptable for practical networks. The *Distributed Gossip* scheme proposed in this section has been designed to overcome the above-mentioned limitation of the Global Gossip scheme, while preserving its main advantages: high probability of gossip coverage, scalability, and simplicity. The only additional assumption imposed is that each node knows the number of its one-hop neighbors, which can be achieved by a periodic beckoning process.

In *Distributed Gossip*, the gossip initiator broadcasts the gossip to all its immediate neighbors. Any node n , other than the initiator, rebroadcasts the gossip upon receiving it for the first time with the following probability:

$$p_d(n) = \min(1, \theta^{-1}(\gamma)/deg(n)), \quad (5)$$

where $deg(n)$ is the number of neighbors of the rebroadcasting node n , and γ is the minimal required *coverage level*. Thus, if a node has no more than $\lceil \theta^{-1}(\gamma) \rceil$ neighbors, it gossips with probability 1; while node n with $deg(n) > \lceil \theta^{-1}(\gamma) \rceil$ gossips with probability $\frac{\theta^{-1}(\gamma)}{deg(n)}$. Distributed Gossip is a localized version of the Global Gossip scheme; instead of using the average node degree, the nodes use their own local degree to set their gossiping probability.

B. Analytical Model

In analyzing the *Distributed Gossip*, we follow the same *a priori decision* method as in Section IV. We assume that before the gossip spread starts, each node decides independently which of the two sets of nodes, active or passive, it belongs to. However, the *Continuum Percolation Theory* cannot be applied anymore, because the active nodes are no longer distributed according to a homogeneous Poisson process. Therefore, we consider a slightly generalized version of the Distributed Gossip scheme, where a node n decides to gossip (or becomes active) with probability

$$p_d(n) = \min(1, c/deg(n)) \quad (6)$$

for some $c > 0$.

Theorem 2: For any $0 < \gamma < 1$, there exist ϕ' and c' such that for any $\phi > \phi'$ and for any $c > c'$, if the Distributed Gossip algorithm is executed according to equation (6) in networks of the average node degree ϕ , the following holds true (proof is omitted due to space limitations):

- (a) a gossip run is successful with probability $\gamma'(\phi, c) \geq \gamma$;
- (b) in a successful gossip run, an arbitrary node receives the gossip message with probability $\gamma'(\phi, c) \geq \gamma$.

The proof of Theorem 2 consists of multiple elaborate stages and is analogous to the analysis done in [15] for NAPs, a randomized sleeping schedule algorithm. Even though the purpose and the design assumptions of the NAP scheme and the Distributed Gossip algorithm are different, the analysis done in [15] can be reused in the Distributed Gossip case by proving the statement of the Lemma 6 in [15] based on the assumptions specific to the Distributed Gossip scheme.

Having shown the existence of the sufficient condition for the Distributed Gossip being able to provide required coverage level γ (when nodes gossip according to equation (6)), next we study how the actual Distributed Gossip, which operates according to equation (5), fares when compared to Global Gossip.

First, we compare the participation of nodes in gossiping in the Distributed and the Global schemes. We introduce the parameter $\delta(\phi)$, which is defined as the ratio between the probability that a node belongs to the *active* set S_a^d in the Distributed Gossip scheme and the probability that a node belongs to the active set S_a in the Global Gossip scheme, assuming the average node degree of $\phi \geq \theta^{-1}(\gamma)$ and some coverage level γ :

$$\delta(\phi) = \frac{P(n \in S_a^d)}{P(n \in S_a)} = \frac{\phi}{\theta^{-1}(\gamma)} P(\text{deg}(n) < \lceil \theta^{-1}(\gamma) \rceil) + \left(P(\text{deg}(n) > \lceil \theta^{-1}(\gamma) \rceil) + \sum_{k=\lceil \theta^{-1}(\gamma) \rceil}^{\infty} \frac{e^{-\phi} \phi^{k+1}}{k(k+1)!} \right).$$

Once $\delta(\phi)$ exceeds 1 (we term this point ϕ_u), it reaches a single maximum and converges to 1 from above as $\phi \rightarrow \infty$. Hence, for $\phi > \phi_u$, the fraction of active nodes in the Distributed Gossip scheme is greater than that in the Global Gossip scheme for networks with the same average node degree.

Second, the node's decision of whether to join the active set or not is more *intelligent* in the Distributed Gossip case. If there are *too few* neighbors ($\text{deg}(n) < \phi$), the gossiping probability of the node running the distributed scheme is greater than that of the same node running the global scheme, lowering the chances of the gossip to die off. On the other hand, the nodes with *too many* neighbors ($\text{deg}(n) > \phi$) would transmit with smaller probability in the distributed case than in the global case, thus lowering the transmission redundancy.

These observations suggest that in the Distributed Gossip scheme, a gossip is more likely to spread to infinitely many nodes than in the Global Gossip case. We formalize our intuition in the following conjecture.

TABLE I
SUMMARY OF 1000 RUNS OF GLOBAL GOSSIP FOR EACH γ

	$\gamma = .9$	$\gamma = .95$	$\gamma = .99$
% of successful runs	90.4	93.8	99.5
% nodes receiving gossip in succ. run	89.83	95.18	98.97

Conjecture 3: When the Distributed Gossip scheme is employed according to equation (5) in networks with average node degree $\phi > \phi_u$,

- (a) a gossip run is successful with probability $\gamma_d(\phi) \geq \gamma$;
- (b) in a successful gossip run, an arbitrary node receives the gossip message with probability $\gamma_d(\phi) \geq \gamma$.

C. Scalability of the Distributed Gossip

For a network of N nodes distributed uniformly and randomly over an area A , the expected number of transmissions in a successful run of the Distributed Gossip scheme is:

$$E[N_{tx}^d] = \gamma_d(\phi) P(n \in S_a^d) N \leq \theta^{-1}(\gamma) \frac{A}{\pi r^2} \max_{\phi \geq \theta^{-1}(\gamma)} \delta(\phi). \quad (7)$$

Thus, the expected number of transmissions in a successful run of the *Distributed Gossip* is upper-bounded by a value that does not depend on the number of nodes in the network, but rather on the area of the network.

VI. SIMULATIONS AND PERFORMANCE

We carried our simulations in SWANS [16], a wireless network simulator for very large networks. The different values of ϕ were obtained by varying the number of nodes in square-shaped networks of the size parameter $\alpha = 40$. The ideal MAC layer was used to concentrate on the fundamental properties of gossiping. In all the simulation runs, the gossip initiator was chosen randomly among the network nodes. In all our experiments, we considered a gossip run successful if the fraction of nodes receiving the gossip was at least 1/2.

A. Global Gossip performance

Table I provides a summary of a larger experiment with 1000 Global Gossip in randomly generated networks of average node degree $\phi = 11.3$ for different values of γ . This confirms one of the main results of this paper – *Theorem 1*.

Fig. 3 displays the Global Gossip simulation results for networks of fixed size and various average node degrees. We observe that, as long as $\phi > \theta^{-1}(\gamma)$, the fraction of nodes covered in a successful gossip run corresponds to the required gossip coverage γ and that the average number of transmissions in a successful gossip run is approximately the expected constant.

B. Distributed Gossip performance

Table II summaries the outcome of 1000 runs of the Distributed Gossip for various values of γ . Clearly, the Distributed Gossip performance exceeds the pre-set value of γ , as it was concluded in Conjecture 3. The price of this “too good” Distributed Gossip performance is a slightly increased number of transmissions in a successful gossip spread, as we can see

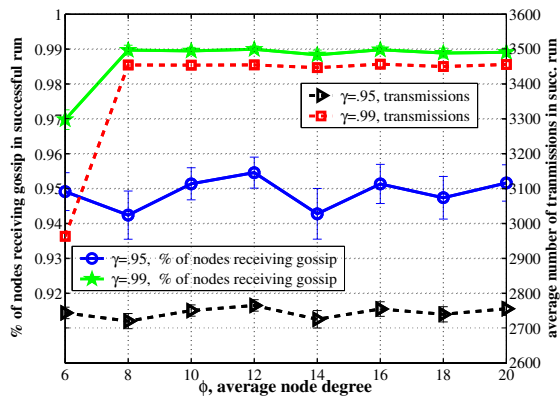


Fig. 3. Global Gossip performance for different γ and ϕ

TABLE II
SUMMARY OF 1000 RUNS OF DISTRIBUTED GOSSIP FOR EACH γ

	$\gamma = .9$	$\gamma = .95$	$\gamma = .99$
% of successful runs	99.0	99.1	99.7
% nodes receiving gossip in succ. run	98.17	99.2	99.87

by comparing the average number of transmissions between Fig. 3 and Fig. 4. However, we should emphasize that the Distributed Gossip achieves at least the pre-set gossip coverage γ in a completely decentralized manner for any average node degree $\phi > \phi_u$. As observed in Fig. 4, the average number of transmissions in both cases can slightly exceed the corresponding upper-bound as per equation (7). This happens due to the finiteness of the simulation area; the nodes located close to the boundary unnecessarily inflate their gossiping probability since they have fewer neighbors.

VII. SUMMARY

In this paper, we developed an analytical model of the Global Gossip scheme, which is a simple probabilistic network-wide broadcasting algorithm. In our model, any gossiping coverage level can be achieved by setting the gossip probability in accordance with the average node degree. We demonstrated that Global Gossip is scalable; that is, the expected number of gossip transmissions in a successful gossip run grows at most linearly with the area of the network and, in particular, is independent of the number of nodes in the network.

Next, we introduced the Distributed Gossip scheme, in which each node sets its own gossiping probability locally, based on the number of its neighbors. We found the sufficient condition for the scheme to achieve the desired performance. We conjectured that, given the same network parameters, the Distributed Gossip scheme has higher probability of a successful gossip spread than the Global Gossip scheme. Moreover, the expected coverage in a successful Distributed Gossip run is larger than that in the Global Gossip.

Our simulations confirmed the key properties of the schemes: they achieve or exceed the pre-set required gossip coverage, and they are scalable with respect to the number of nodes in the network. Although, in principle, both of

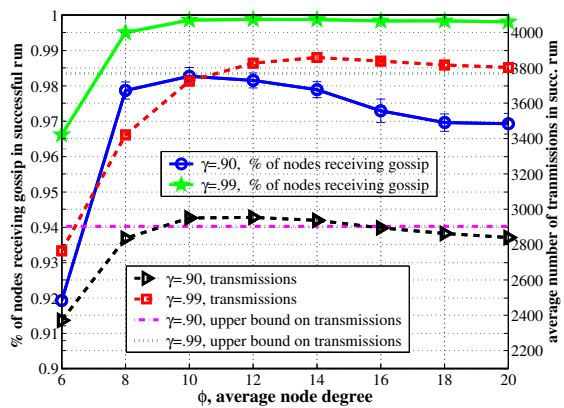


Fig. 4. Distributed Gossip performance for different γ and ϕ

the gossiping schemes could be applied to networks with non-uniform spatial distribution of nodes, nevertheless, the Distributed Gossip scheme will perform better due to its local adaptability to spatial node density variations.

VIII. ACKNOWLEDGEMENT

The work of Zygmunt J. Haas was supported in part by the NSF grants ANI-0329905 and CNS-0626751 and by the AFOSR contract FA9550-09-1-0121.

REFERENCES

- [1] Charles E. Perkins, *Ad Hoc Networking*, Addison Wesley, 2001.
- [2] Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen, and Jang-Ping Sheu, "The broadcast storm problem in a mobile ad hoc network," in *Proc. ACM MobiCom*, August 15-20, 1999, pp. 151–162.
- [3] Zygmunt J. Haas, Joseph Y. Halpern, and Erran L. Li, "Gossip-based ad hoc routing," *IEEE Trans. on Networking*, vol. 14, pp. 479–491, June 2006.
- [4] Yoav Sasson, David Cavin, and André Schiper, "Probabilistic broadcast for flooding in wireless mobile ad hoc networks," in *IEEE WCNC 2003*, March 16-20, 2003.
- [5] Mathew Penrose, *Random Geometric Graphs*, Oxford Univ. Press, 2003.
- [6] Brad Williams and Tracy Camp, "Comparison of broadcasting techniques for mobile ad hoc networks," in *Proc. ACM MobiHoc*, June 9-11, 2002, pp. 194–205.
- [7] I. Stojmenovic and J. Wu, "Broadcasting and activity-scheduling in ad hoc networks," in *Ad Hoc Networking*. IEEE Press, 2004.
- [8] Julien Cartigny and David Simplot, "Border node retransmission based probabilistic broadcast protocols in ad-hoc networks," in *HICSS '03: Proc. of the 36th Annual Hawaii International Conference on System Sciences*, Washington, DC, USA, 2003, p. 303.
- [9] E. N. Gilbert, "Random plane networks," *J. Soc. Indust. Appl. Math.*, vol. 9, no. 4, pp. 533–543, December 1961.
- [10] Piyush Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, Edited by W.M. McEneaney, G. Yin, and Q. Zhang, pp. 547–566, 1998.
- [11] Olivier Dousse, Patrick Thiran, and Martin Hasler, "Connectivity in ad-hoc and hybrid networks," in *IEEE INFOCOM 2002*, June 2002.
- [12] Ronald Meester and Rahul Roy, *Continuum Percolation*, Cambridge University Press, 1996.
- [13] S. Torquato, *Random Heterogeneous Material, Microstructure and Macroscopic Properties, Mechanics and Materials*, Springer, 2002.
- [14] Rajeev Motwani and Prabhakar Raghavan, *Randomized Algorithms*, Cambridge University Press, 1995.
- [15] P. Brighten Godfrey and David Ratajczak, "Naps: scalable, robust topology management in wireless ad hoc networks," in *ISPN 2004*, April 26-27, 2004, pp. 443–451.
- [16] Rimon Barr and Zygmunt J. Haas, JIST/SWANS, <http://jist.ece.cornell.edu/>, 2004.