A fair comparison of gossip algorithms over large-scale random topologies

Ruijing Hu*, Julien Sopena*, Luciana Arantes*, Pierre Sens*, and Isabelle Demeure†
*Université Pierre et Marie Curie, CNRS/INRIA, Paris, France
†Institut Telecom, Telecom ParisTech, CNRS, Paris, France

***Abstract*—This paper presents a thorough performance evaluation comparison study of three widespread used gossip probabilistic algorithms over well-known random graphs that represent some large-scale network topologies: Bernoulli (or Erdős-Rényi) graph, random geometric graph, and scale-free graph. In order to conduct such a comparison fairly, notably in terms of reliability, we propose a new parameter, denoted *effectual fanout*. For a given topology and gossip algorithm, the effectual fanout characterizes the mean dissemination power of infected sites. For large-scale networks, the effectual fanout has thus a strong linear correlation with message complexity. It enables the accurate analysis of the behavior of a gossip algorithm over a topology while also simplifies the theoretical comparison of different gossip algorithms on this topology. Based on extensive comparison simulation experiments, that use the effectual fanout, on top of OMNet++, we discuss the impact of topologies and gossip algorithms on performance, and of how to combine them to have the best gain in terms of reliability.**

I. Introduction

Information dissemination, in which a site attempts to broadcast messages to all the other sites of the network, is essential for many distributed systems and applications, including large-scale ones. On the other hand, the latter usually require a dissemination protocol that provides high reliability, which expresses the percentage of broadcast messages that are received by all sites of the system, with both low latency and message complexity.

A straightforward but inefficient way to disseminate information network wide is pure flooding protocol in which upon the first reception of a message, every site of the network relays it once to its respective neighbors [18]. However, in this case, a very large number of messages may be generated, which entails broadcast storm problems [23]. To mitigate this undesirable phenomenon, probabilistic gossip algorithms have emerged as a solution to implement effective broadcast protocols, highlighted by their simplicity, high reliability, and scalability [13], [30]. Either applied to overlay networks [10], [12], [19], or exploited in wireless ad hoc and sensor networks [4], [15], [17], [29], [31], they reduce the number of messages and well satisfy application constraints. Nevertheless, probabilistic gossip protocols do not always ensure 100% of reliability. Hence, aiming a tradeoff between reliability and message complexity, algorithms found in the literature have input parameters of different natures. For instance, fixed fanout gossip (*GossipFF*) [10], [12], [19], probabilistic edge gossip (*GossipPE*) [15], [29], and probabilistic broadcast gossip (*GossipPB*) [4], [17], [31] are well-known gossip algorithms: *GossipFF* applies as input the fanout, which is the number of per site target neighbors to send the message; in *GossipPE*, based on an input probability parameter, a site randomly chooses those edges over which received message should be retransmitted; in *GossipPB*, the input parameter defines the probability that a site broadcasts the message to all its neighbors.

Besides gossip configuration parameters, network topologies have their own proprieties (e.g., the degree distribution, the edge dependency, etc.), which also have an impact in the performance of gossip algorithms such as message complexity, reliability, and latency. Thus, in this paper, we consider the following widespread studied random graphs: Bernoulli (or Erdős-Rényi) graph $\mathcal{B}(N, p_N)$ [8], random geometric graph $\mathcal{G}(N, \rho)$ [25], and scale-free graph $\mathcal{S}(N, m)$ [2] which respectively model peer-to-peer system [19], wireless sensor network [17], and ad hoc network [15]

Considering the above discussed differences, we propose in this work to compare the performance of gossip probabilistic algorithms on different topologies. However, in order to carry out a fair comparison, we have introduced a new parameter, denoted *effectual fanout* which expresses the average number of messages per retransmission. It characterizes, therefore, the potential mean dissemination power of infected sites, i.e., those that received at least once the message. In large-scale systems, the effectual fanout has thus a strong linear correlation with message complexity metric as we show and prove in Section V. For a given value, the *effectual fanout* can be analytically calculated in function of the input parameter of the corresponding gossip algorithm (e.g., fanout, probability, etc.) and thus it simplifies the theoretical comparison of different gossip algorithms on a fixed topology. The advantage of using the effectual time compared to message complexity metric is that the former can be easily calculated analytically while the latter requires to know the total number of messages generated by each algorithm in function of the topology.

Exploiting the *effectual fanout* parameter, we present in this paper results of an extensive performance evaluation, conducted on top of OMNET++ [1], which compare *GossipFF*, *GossipPE*, and *GossipPB* algorithms over the above mentioned three topologies. To the best of our knowledge, it is the first time that such a comparison study has been proposed.

The remainder of this paper is organized as follows. Section II and III give an overview of our system random networks and probabilistic gossip algorithms. Section IV introduces the performance metrics. The effectual fanout is presented in Section V. Section VI shows simulation results on OMNET++ while Section VII discusses some related work. Finally, Section VIII concludes this work.

II. System Topologies

In the sequel, $\mid l \mid$ denotes the size of set l.

We consider a large-scale dissemination system Π comprised of N sites $\{s_1, s_2, \cdots, s_N\}$. The set of all s_i's neighbors is denoted Λ_i and $V_i = \mid \Lambda_i \mid$ denotes the degree of s_i; $P(k)$ represents the degree distribution of a site with k neighbors (i.e., the fraction of sites with degree k) in the graph and $\bar{V}$ is the mean degree. Therefore, $\bar{V} = \sum_{k=0}^{N-1} P(k) \cdot k$. There is no message loss.

Three random topologies are taken into account in our study: Bernoulli (or Erdős-Rényi) graph $\mathcal{B}(N, p_N)$ [8], random geometric graph $\mathcal{G}(N, \rho)$ [25], and scale-free graph $\mathcal{S}(N, m)$ [2]. See Figures 1(a), 1(b), and 1(c) respectively.

Bernoulli (or Erdős-Rényi) graph $\mathcal{B}(N, p_N)$ is a random bidirectional graph constructed by connecting sites randomly with probability p_N, independently of other edges. Based on [9], we suppose that $p_N > \frac{(1+\varepsilon)\cdot\ln(N)}{N}$, with a positive constant ε, aiming at having a *giant component* which would have N sites with Poisson-law degree distribution $P(k) = \exp\left(-\bar{V}\right) \frac{\bar{V}^k}{k!}$, where $\bar{V} = p_N \cdot N$.

The random geometric graph $\mathcal{G}(N, \rho)$ is a random bidirectional graph drawn on a bounded region. In this article, such a region is a rectangular plane with length a and width b. $\mathcal{G}(N, \rho)$ is generated by placing sites uniformly at random and independently on the region. Furthermore, two sites are connected, whenever the distance between them is at most ρ. Moreover, based on [26], we can fine-tune $\rho > \sqrt{\frac{(1+\varepsilon)\cdot\ln(N)\cdot a\cdot b}{N\cdot\pi}}$ with a positive constant ε in order to ensure that the graph is connected with Poisson-law degree distribution [20] such that $P(k) = \exp\left(-\bar{V}\right) \frac{\bar{V}^k}{k!}$, where $\bar{V} = \frac{N\cdot\pi\cdot\rho^2}{a\cdot b} - 1$ when ignoring the border effect of the region in $\mathcal{G}(N, \rho)$.

Scale-free graph $\mathcal{S}(N, m)$ is a random bidirectional graph generated by Barabási-Albert model [2]. Starting from a small $clique$ of m_0 sites, at every time step a new site is added such that its m ($\leqslant m_0 \ll N$) edges connect it to m different sites already present in the graph. The probability p that a new site will be connected to an existed site is proportional to the degree of the latter. This is called preferential attachment. This process ensures that the graph is connected with power-law degree distribution approximately equal to $P(k) = \frac{2m(m+1)}{k(k+1)(k+2)}$ where $k = m, m+1, \cdots, N-1$ and $\bar{V} = 2m$ which does not depend on N [27]. In this network, there are **hub** and **periphery** sites which have degree greater than $2m$ and between m and $2m$ respectively. Hence, the system Π is composed by the set of hubs denoted Π_h and the set of peripheries denoted Π_p. We can deduce that $\mid \Pi_p \mid > 3 \mid \Pi_h \mid$.

Definition 1. Edge Dependency (or Clustering Coefficient) *of a given random graph, for distinct sites* s_i,s_j,s_k, *is defined as the conditional probability that, given the existence of edges* $s_i \sim s_k$ *and* $s_j \sim s_k$, *an edge* $s_i \sim s_j$ *also exists (i.e.,* $P\left(s_i \sim s_j | s_i \sim s_k, s_j \sim s_k\right)$*).*

In [3], it has been proved that $\mathcal{B}(N, p_N)$ has little edge dependency, i.e., the existence of an edge over $\mathcal{B}(N, p_N)$ does not depend on the others. Therefore, $P\left(s_i \sim s_j | s_i \sim s_k, s_j \sim s_k\right) = P\left(s_i \sim s_j\right) = p_N$. On the contrary, $\mathcal{G}(N, \rho)$ presents a high edge dependency and the existence of edges is correlated. More precisely, when border effect is neglected $P\left(s_i \sim s_j | s_i \sim s_k, s_j \sim s_k\right) = 0.5865$, a value greater than the probability p_N in $\mathcal{B}(N, p_N)$.

Notice that, contrarily to $\mathcal{G}(N, \rho)$, the site that gossips a message over $\mathcal{B}(N, p_N)$ may independently send it to any other neighboring sites or receive it from them.

In [7], it is shown that edge dependency over $\mathcal{S}(N, m)$ is very low, i.e., in the same order of $\mathcal{B}(N, p_N)$. On the other hand, unlike $\mathcal{B}(N, p_N)$ and $\mathcal{G}(N, \rho)$, the degree variance is quite high, due to its power distribution.

Compared to the other two graphs, $\mathcal{S}(N, m)$ has the smallest diameter due to the *hubs* that create short-cut paths [5]. The diameter of $\mathcal{B}(N, p_N)$ is small with rare cliques whereas in $\mathcal{G}(N, \rho)$, the diameter tends to be large and many small cliques turn out.

III. Gossip Algorithms

Information dissemination in large-scale network is commonly studied on basis of Algorithm 1. Initially, the source sends a message to *all* of its neighbors (lines 2 and 3). A site delivers and retransmits a received message provided it has not previously received it; otherwise the message is discarded. The sites that have received at least once the message are denoted **infected sites** while those that received no message are denoted **isolated sites**.

Algorithm 1: Generic Gossip algorithm

```
1 Broadcast (⟨msg⟩)
2     foreach s_j ∈ Λ_i do
3         Send(⟨msg⟩, s_j)

4 Receive (⟨msg⟩)
5     if msg ∉ msgHistory then
6         Deliver(⟨msg⟩) ;
7         msgHistory ← msgHistory ∪ {⟨msg⟩} ;
8         Gossip(⟨msg⟩,parameters) ;
```

There are three main probabilistic gossip families to implement the retransmission Gossip() procedure, namely (1) Fixe Fanout gossip (*GossipFF*), (2) Probabilistic Edge gossip (*GossipPE*), and (3) Probabilistic Broadcast gossip (*GossipPB*). Besides the received message, all these algorithms receive one or more parameters whose value is the same for all sites.

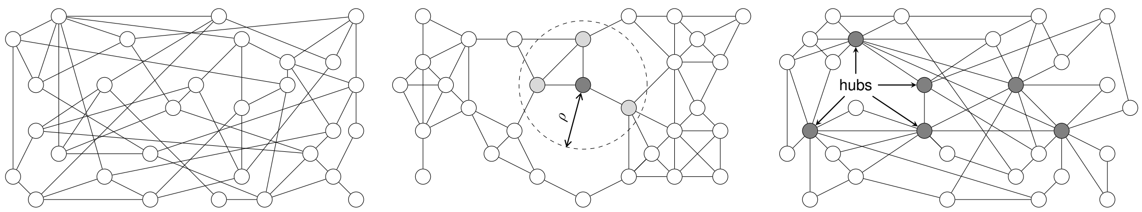

(a) Bernoulli Graph (b) Random Geometric Graph (c) Scale-Free Graph

Figure 1. Examples of the three random topologies with 30 sites and mean degree=4

9 /* $fanout$: number of selected neighbors */
10 **GossipFF** (⟨*msg*⟩,*fanout*)
11 **if** $fanout \geqslant V_i$ **then**
12 $toSend \leftarrow \Lambda_i$
13 **else**
14 $toSend \leftarrow \emptyset$
15 **for** $f = 1$ *to* $fanout$ **do**
16 random select $s_j \in \Lambda_i / toSend$
17 $toSend \leftarrow toSend \bigcup s_j$
18 **foreach** $s_j \in toSend$ **do**
19 Send(⟨*msg*⟩, s_j)

Algorithm 2: Fixe Fanout Gossip (at s_i)

In *GossipFF* (Algorithm 2), site s_i sends msg to a fixed number of sites, denoted $fanout$, in Λ_i, which are randomly selected (lines 15-17). Notice that if $fanout \geqslant V_i$, s_i transmits msg to all its neighbors (lines 11 and 12). Particularly, if $fanout \geqslant \max\{V_1, V_2, \cdots V_N\}$, Algorithm 2 is a pure flooding algorithm.

20 /* p_e: probability to use an edge */
21 ***GossipPE*** (⟨*msg*⟩,p_e)
22 **foreach** $s_j \in \Lambda_i$ **do**
23 **if** `Random()` $\leqslant p_e$ **then**
24 Send(⟨*msg*⟩, s_j)

Algorithm 3: Probabilistic Edge Gossip (at s_i)

For the two following algorithms, `Random()` generates a random number in the interval $[0, 1]$.

In *GossipPE* (Algorithm 3), site s_i randomly chooses those edges over which msg should be transmitted with regard to a fixed probability p_e (see line 23). Note that when $p_e = 1$ for all sites, we obtain the flooding algorithm.

25 /* p_v: probability to broadcast */
26 ***GossipPB*** (⟨*msg*⟩,p_v)
27 **if** `Random()` $\leqslant p_v$ **then**
28 **foreach** $s_j \in \Lambda_i$ **do**
29 Send(⟨*msg*⟩, s_j)

Algorithm 4: Probabilistic Broadcast Gossip (at s_i)

Unlike Algorithm 3, in *GossipPB* (Algorithm 4), each site, except the source, diffuses msg to all its neighbors with fixe probability p_v (see line 27). In particular, when $p_v = 1$ this protocol becomes the flooding algorithm.

IV. Performance Metrics

In the context of information dissemination, the following metrics are used in the literature [17], [19], [21] for performance evaluation:

Message Complexity, denoted $\mathbf{M}$**:** measures the mean number of messages received (or sent, since no message loss is taken into account) by each site:

$$M = \frac{\Omega}{N-1} \tag{1}$$

where Ω is the total number of messages exchanged during the dissemination.

Fraction of Total Infected Sites, denoted α**:** is defined as the percentage of all sites in the system that delivered a message generated by a source in the end of the dissemination.

Reliability, denoted R: is defined as the percentage of messages generated by a source that are delivered by all sites. A reliability value of 100% is indicative that the algorithm was successful in delivering any given message to all sites (i.e., $\alpha = 100\%$ for any given message) ensuring thus *atomicity* similarly to pure flooding algorithms [19].

Latency, denoted L: measures the number of hops required to deliver a message to all recipients, i.e., the number of hops of the longest path among all the shortest paths from the source to all other sites that received the message.

An efficient dissemination algorithm aims at providing both large fraction of total infected sites and high reliability, while minimizing both message complexity and latency.

V. Effectual Fanout

The number of retransmitted messages of the three gossip algorithms, and therefore their message complexity, depends on their respective input parameters (p_v, p_e or $fanout$), which are, in fact, quite different. Hence, aiming at conducting a fair uniform comparison of these algorithms over the topologies described in Section II, we have introduce a new parameter denoted ***effectual fanout*** F_{eff}. The latter enables the accurate analysis of the behavior of a gossip algorithm over a topology while also simplifies the theoretical comparison of different gossip algorithms on this topology. For a fixed topology and gossip algorithm, the effectual fanout characterizes the mean dissemination power of infected sites.

Therefore, when the number of sites of the system is very large, the effectual fanout has a strong linear correlation with message complexity, as shown in Theorem 2 of the current section. Notice that in function of both an algorithm and a topology, it is possible, for a given effectual fanout, to deduce the value of the mentioned input parameter of the gossip algorithm in question, as shown in the following.

Based on *GossipPE*, *GossipPB*, and *GossipFF* algorithms, we define respectively that:

$$F_{eff}^{\mathit{GossipPE}} = p_e \cdot \bar{V} \tag{2}$$

$$F_{eff}^{\mathit{GossipPB}} = p_v \cdot \bar{V} \tag{3}$$

$$F_{eff}^{\mathit{GossipFF}} = \sum_{k=1}^{fanout-1} P(k) \cdot k + \sum_{k=fanout}^{N-1} P(k) \cdot fanout \tag{4}$$

due to the two conditions in Algorithm 2 (lines 11 and 13).

We respectively denote $\mathbf{U_h}$ and $\mathbf{I_h}$ the expected number of sites that have not been infected before the end of hop h and the the expected number of newly infected sites within hop h, for $1 \leqslant h \leqslant L$ where L is the latency. Observe that U_0 equals to $N-1$, I_0 equals to 1, and $U_L = (1-\alpha)N$.

For hop h, U_h and I_h are related as follows:

$$I_h = U_{h-1} - U_h,\ 1 \leqslant h \leqslant L \tag{5}$$

Theorem 2. *For the three probabilistic gossip algorithms over the three large-scale random topologies ($N \gg 1$), the message complexity $M \approx \alpha F_{eff}$.*

Proof: Since there is no loss of messages, the total number of messages received by each site is equal to the number of transmitted messages. In every hop h, a site will relay F_{eff} messages to its neighbors, while the expected number of newly infected sites in hop h is I_h. Thus, the expected number of transmitted messages in hop h is $F_{eff} \cdot I_h$.

Considering all hops and (5), we obtain Ω, the total number of received messages:

$$\Omega = \sum_{h=1}^{L} F_{eff} \cdot I_h = F_{eff} \cdot \sum_{h=1}^{L} I_h = F_{eff} \cdot (N - 1 - U_L)$$

By Equation (1), $M = \frac{(N-1-U_L)}{N-1} \cdot F_{eff} = \frac{(\alpha N-1)}{N-1} \cdot F_{eff}$, since N is very large we then get $M = \alpha F_{eff}$. ■

Corollary 3. *In order to have high reliability for the three probabilistic gossip algorithms over the three large-scale random topologies, message complexity $M \approx F_{eff}$.*

Proof: When the high reliability is reached (e.g., heuristically, over 95% of *total* sites are infected on average at the end of a message dissemination), α is very close to 100% and, according to Theorem 2, the result is obtained. ■

VI. Performance Evaluation

In this section, based on simulation experiments conducted on top of OMNET++ [1], we evaluate the performance of the three algorithms described in Section III over each of the large-scale random topologies described in Section II based on the metrics presented in Section IV.

Our goal is to fairly compare the three algorithms over the three topologies. To this end, in the experiments, the value of the respective input parameter of each algorithm has been varied and, in order to unify them, the effectual fanout for the different values have been computed (see section V). The performance evaluation results are then presented in function of the effectual funout: for different values of the effectual fanout, we evaluate the percentage of infected sites, reliability, and latency of the three gossip algorithms over the three topologies.

For each gossip algorithm, the mean of 200 different messages are generated by 200 different sources that are chosen uniformly amongst 1000 sites over 20 different graphs related to each of the topologies. Their comparisons are based on effectual fanout for the 200 message disseminations.

We consider that he network is composed of $N = 1000$ sites and, in order to ensure connectivity, $\varepsilon = 1$ for $\mathcal{B}(N, p_N)$ and $\mathcal{G}(N, \rho)$. Since we aim at having almost the same mean degree for all topologies ($\bar{V} \approx 14.0$), the following topology parameters were chosen, as shown in Table I:

TOPOLOGY	PARAMETERS
$\mathcal{B}(N, p_N)$	$p_N = 0.014$
$\mathcal{G}(N, \rho)$	$a = 7500$, $b = 3000$, $\rho = 330$
$\mathcal{S}(N, m)$	$m_0 = 9$ ($m_0 - clique$), $m = 7$

Table I
Topology Parameters

A. Linearity Between Effectual Fanout and Message Complexity

As explained, our aim is to fairly compare the performance of the three gossip algorithms, notably the reliability, using the effectual fanout. Such a fairness requires the equivalence in terms of message complexity of the three algorithms over a given topology. Therefore, we would like to verify the linear relation between the effectual fanout and message complexity. Figures 2(a), 2(b), and 2(c) show this relation. They confirm that the linearity $F_{eff} = M$ holds whenever F_{eff} value is great enough ($F_{eff} > 3$). On the other hand, for smaller F_{eff} values, the fraction of infected sites α is too small to be neglected in the equation of Theorem 2. The only exception is for *GossipPB* over $\mathcal{G}(N, \rho)$ (Figure 2(b)) since, in this case, a clustering effect (see Section VI-C) prevents this algorithm to benefit from the growth of the dissemination power. Notice that such value is much smaller than the dissemination power value which provides reliability. Hence, the fairness of the algorithm comparison is ensured in this case.

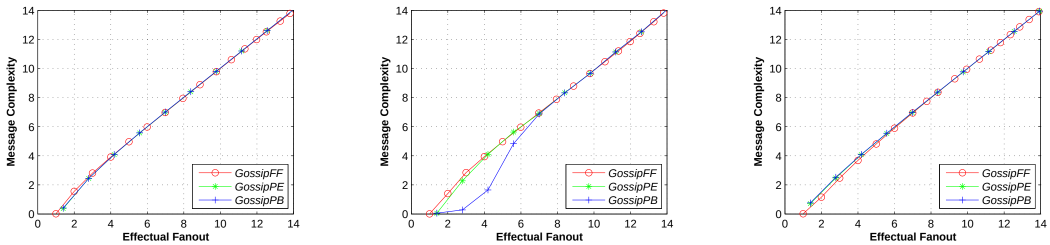

(a) Over $\mathcal{B}(N, p_N)$ (b) Over $\mathcal{G}(N, \rho)$ (c) Over $\mathcal{S}(N, m)$

Figure 2. Relation between Message Complexity and Effectual Fanout

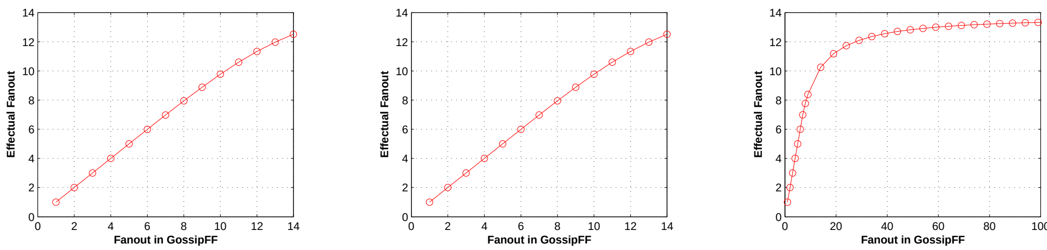

(a) Over $\mathcal{B}(N, p_N)$ (b) Over $\mathcal{G}(N, \rho)$ (c) Over $\mathcal{S}(N, m)$

Figure 3. Difference between Fanout in *GossipFF* and Effectual Fanout.

Fanout vs. Effectual Fanout: In order to well understand the difference between the fanout in *GossipFF* and effectual fanout, Figures 3(a), 3(b), and 3(c) present the value of effectual fanout for each random topology in function of fanout in *GossipFF*. We observe that for small values there is an equality between them, while their values diverge for large fanout values (i.e., the effectual fanout is proportionally smaller). In fact, the sites whose number of neighbors is inferior to fanout do not use all their dissemination power. In $\mathcal{S}(N, m)$, where the degree variance is very great and the number of sites such as peripheries with small number of neighbors is very large, this phenomenon is much more remarkable (see Figure 3(c)). Hence, this shows the importance of the hypotheses done in a great number of theoretical studies to have a fanout for *GossipFF* inferior to the minimum degree of the graph.

B. Algorithms Equivalence Over $\mathcal{B}(N, p_N)$

We discuss now the performance evaluation for the gossip algorithms over $\mathcal{B}(N, p_N)$. On one hand, both the fraction of infected sites (α) and the reliability (R) in Figures 4(a) and 4(b) present a threshold effect as a function of effectual fanout. In other words, the fraction of infected sites or the reliability stays at 0 for some small effectual fanout values, but it quickly comes to 100% for a threshold value (effectual fanout = 4). On the other hand, we observe that the performance for all gossip algorithms is the same for the same effectual fanout.

However, if we compare the thresholds for the fraction of infected sites and reliability, they are different. The fraction of infected sites percolates with smaller effectual fanout than reliability (respectively, 4 and 13). As a matter of fact when the effectual fanout is great enough such that almost all sites receive each message (i.e., $\alpha \approx 100\%$), none of the messages is received by all sites (i.e., $R = 0$). Only when the effectual fanout equals to 13 that almost all messages are received by every site surely (i.e., $R \approx 100\%$). We thus observe a great gap in terms of effectual fanout value between the dissemination power necessary for infecting almost every site and high reliability.

Since the algorithms have the same behavior over $\mathcal{B}(N, p_N)$, then, we can use the theoretical result of *GossipFF* [19] to determine the corresponding thresholds for *GossipPE* and *GossipPB*: $fanout = -\ln\left(\frac{-\ln(R)}{N}\right)$. For instance, for $R = 99.4\%$, $fanout = -\ln\left(\frac{-\ln(.994)}{1000}\right) \approx 12$. By Equation (4) in Section V we obtain $F_{eff} = 11.3$. Thereby, $p_e = p_v = F_{eff}/\bar{V} = 11.3/14 = 0.81$. Hence, it becomes possible to dimension the input probabilities of *GossipPE* and *GossipPB* to obtain a desired reliability.

In Figure 4(c), after a given effectual fanout, latency does not decrease anymore, but converges towards pure flooding approach (i.e., the shortest routes between the source and the other sites), and therefore, towards the minimum latency.

C. Algorithms Difference over $\mathcal{G}(N, \rho)$

We now present simulation results related to the performance of the gossip algorithms on $\mathcal{G}(N, \rho)$ (Figures 5 and 6).

If the performance of the gossip algorithms is identical in $\mathcal{B}(N, p_N)$, it is not always the case for other random topologies. Thus, if we now consider the reliability (see

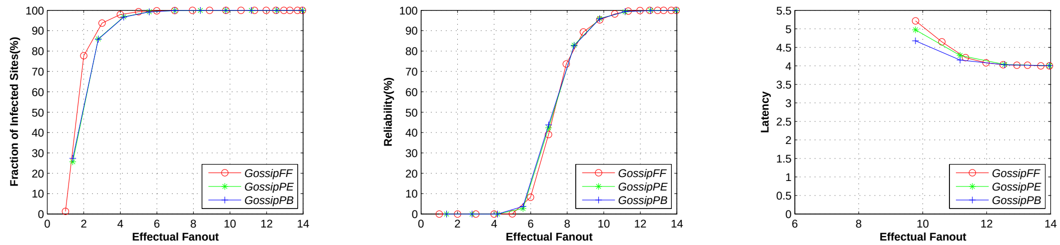

(a) Infected Sites Over $\mathcal{B}(N, p_N)$ (b) Reliability Over $\mathcal{B}(N, p_N)$ (c) Latency Over $\mathcal{B}(N, p_N)$

Figure 4. Performance Comparison of Algorithms over $\mathcal{B}(N, p_N)$.

Figure 5(b)), we notice that *GossipFF* is much more efficient (i.e., with merely $F_{eff} = 8.5$, $R = 99\%$ is reached) than the two other algorithms that require $F_{eff} = 14$ to reach $R = 99\%$. Furthermore, we can observe in the same figure that the threshold effect for both *GossipPE* and *GossipPB* is much smoother (i.e., from 5.5 to 14) than for *GossipFF* and that *GossipPE* presents a slightly better performance than *GossipPB*. However, if we look at the performance in terms of the fraction of infected sites in Figure 5(a), the comparison results are quite different. *GossipPE* has very similar performance to *GossipFF* which is the most effective. *GossipPB* shows the worst performance: it requires about $F_{eff} = 8$ in order to infect almost every sites.

The behavior of the latency curves of Figure 5(c) for $\mathcal{G}(N, \rho)$ is similar to that of $\mathcal{B}(N, p_N)$, except that the minimum latency value is around 22 hops since the diameter of the former is greater than the latter.

In order to thoroughly analyze the results, we conducted a series of specific experimentations when placing the source in the center of rectangular plane of dimension 3000×7000 with 1000 sites uniformly distributed at random. The mean degree corresponding to the radius $\rho = 330$ is about 14. The results are presented in Figure 6. Several values of F_{eff} are chosen for the three gossip algorithms. In addition, the axis *x* and *y* represent the geographic position of the site in the graph, whereas the axis z characterizes the percentage of messages received by every site. The greater the value towards the axis z for a site, the greater the number of messages received by the site. The plane *z=0* indicates the sites that never received any message.

The performance of *GossipFF* is shown in the third column of Figure 6. We can verify that this algorithm is the most effective for infecting all sites (i.e., with F_{eff} merely equal to 6.97) contrarily to the other two algorithms (see Columns 1 and 2) that cannot broadcast every message from the source to the whole system Π until $F_{eff} = 9.77$. Even though these two algorithms completes the broadcast with almost the same performance, the evolution of their infection is quite different.

On the first column, we notice that *GossipPB* presents a peek for long time for several values of F_{eff} in the graphs. This phenomenon implies that the infected sites are located around the source and the message dissemination stops quickly. It can be explained by the clustering effect entailed by the broadcast probability p_v: in this algorithm, sites stop retransmitting the message with probability $1-p_v$. If this probability is high, the sites that do not relay the message give rise to a confinement around the source (i.e., the border of the peak). On the other hand, by increasing p_v, the clustering effect is reduced and the message can be received by every site. The study of such phenomenon is particularly important since, as explained in Section II, $\mathcal{G}(N, \rho)$ has very high edge dependency which induces a higher number of possible peak borders which increase the risk of dissemination stop.

Inversely, for *GossipPE* (see the second column), by slightly increasing the value F_{eff}, almost all sites receive every message from the source. Nevertheless, contrarily to *GossipFF* there are always some sites which receive only a few messages. Such sites are located either on the border of the rectangular plane or in areas with small site density of the graph. As a matter of fact, *GossipPE* imposes random choice for each edge of every site no matter its degree. Therefore, sites having very few neighbors, with high probability, do not receive all messages. This explains *GossipPE* bad reliability (see Figure 5(b)) even when almost all sites are infected (see Figure 5(a)). For instance, when $F_{eff} = 5$, $\alpha \approx 100\%$ whereas there is no reliability.

Therefore, this study shows why *GossipFF* is particularly efficient over $\mathcal{G}(N, \rho)$. By forcing each site to retransmit some messages, it reduces the clustering effect more effectively than the other algorithms. Furthermore, by obliging the sites with small degree (i.e. smaller than $fantout$) to broadcast the message to all its neighbors, it prevents the risk of dissemination stop of small density areas.

D. Algorithms Difference over $\mathcal{S}(N, m)$

We discuss now the performance of the different gossip algorithms over $\mathcal{S}(N, m)$ (Figures 7 and 8). Results are completely different from the other two random topologies. Similarly to $\mathcal{B}(N, p_N)$, the three algorithms present the same performance behavior in terms of the fraction of infected sites as shown in Figure 7(a). Nevertheless, contrarily to the reliability over $\mathcal{G}(N, \rho)$, *GossipFF* turns to be the worst choice (see Figure 7(b)).

Such a performance behavior is a consequence of the

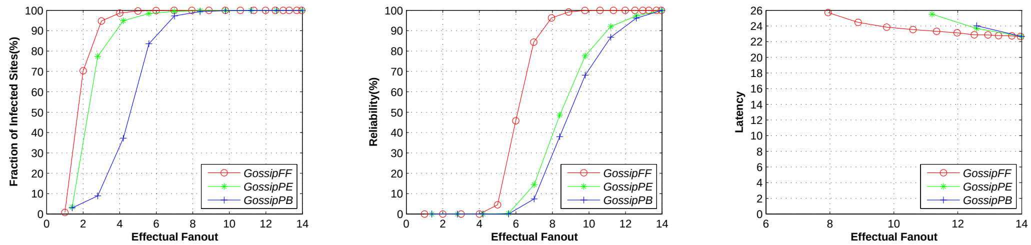

(a) Infected Sites Over $\mathcal{G}(N, \rho)$ (b) Reliability Over $\mathcal{G}(N, \rho)$ (c) Latency Over $\mathcal{G}(N, \rho)$

Figure 5. Performance comparison of algorithms over $\mathcal{G}(N, \rho)$.

degree distribution of the graph which has sites with higher degrees, the **hubs** (see Section II). In order to understand the dissemination power of the hubs, we have measured, for different F_{eff} values, the reliability of the hubs (i.e., the proportion of messages generated by the source that are received by all hubs in the system Π). The results are presented in Figure 8. For each algorithm, the latter are compared with the reliability when we consider all sites of Π (denoted global reliability). This comparison shows that hubs are infected on priority no matter which algorithm is applied. Thereby, with a F_{eff} equal to 6, the reliability of the hubs is 100% for all the three algorithms whereas almost none of the messages is received by all sites (i.e., the global reliability is still zero).

GossipFF presents the worst performance which can be explained by its poor exploitation of hubs. In fact, even if hubs degree is quite high, the algorithm limits their dissemination power to the value fixed by the $fanout$. On the other hand, it should be understood that a transmission of 10 messages by one site is more powerful than a transmission of 1 message by 10 sites. In the first case, all receivers are different, which ensures a better message dissemination with less message redundancy.

The fact that the dissemination potential of hubs is not fully exploited also explains the latency of Figure 7(c). Even if *GossipPB* and *GossipPE* present the same latency, it is not the same for *GossipFF* which keeps a larger latency when its reliability is near 100%. As a matter of fact, in $\mathcal{S}(N, m)$, the hubs is the heart of the network: the peripheries have at least one hub in its neighborhood with high probability. By limiting the dissemination power of the hubs, *GossipFF* discards numerous short-cut paths.

E. Impact of the Topology on the Algorithms

Since in our simulations we have considered topologies with the same mean degree ($\bar{V} \approx 14.0$), we can compare the reliability of the algorithms over the different topologies, as shown in Figure 9 and summarized in Table II. When the graph has low edge dependency and low degree variance as in ($\mathcal{B}(N, p_N)$), the three algorithms present the same behavior. When edge dependency (resp., degree variance) is introduced in the graph, but the degree variance (resp., dependency) does not change as in $\mathcal{G}(N, \rho)$ (resp., $\mathcal{S}(N, m)$), the performance of *GossipPB* (resp., *GossipFF*) decreases.

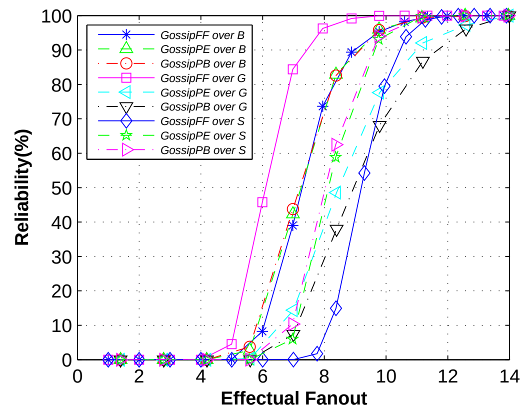

Figure 9. Topologies impact on algorithms

Such results confirm that the best algorithm choice for the reliability with the same message complexity depends on the properties of network topology. It should be pointed out that the performance of *GossipPE* is never worse than *GossipPB*, since in the percolation theory [16], the former can be modeled by bond percolation while the latter matches the site percolation. The percolation threshold of bond percolation is always smaller or equal to that of site percolation in any topology.

	Low Degree Variance	**High Degree Variance**
Low Edge Dependency	$\mathcal{B}(N, p_N)$: ***GossipFF, GossipPE,GossipPB***	$\mathcal{S}(N, m)$: ***GossipPE, GossipPB***
High Edge Dependency	$\mathcal{G}(N, \rho)$: ***GossipFF***	$--$

Table II
ALGORITHM CHOICE

We have also measured the relative gain in terms of effectual fanout when the reliability reaches 80% and 99% and the fraction of infected sites is large (i.e., α is approximate to 100%) compared with effectual fanout needed by the flooding algorithm (i.e., the maximum message complexity). The results over the three random topologies are shown in Tables III and IV respectively. We observe that over $\mathcal{G}(N, \rho)$ to reach $R = 99\%$ the gain of *GossipPB* and *GossipPE* is zero. Hence, they need almost the same message complexity as the pure flooding. On the other hand, to reach $R = 80\%$, $\mathcal{B}(N, p_N)$ exhibits the best gain for *GossipPB* and *GossipPE* amongst all random topologies. Furthermore, *GossipFF* over $\mathcal{G}(N, \rho)$ is the best combination for achieving the highest gain.

In conclusion, in order to reduce message complexity of gossip algorithms in contrast to the flooding algorithm, it is necessary to consider both the gossip algorithm and the topology.

	$\mathcal{B}(N, p_N)$	$\mathcal{S}(N, m)$	$\mathcal{G}(N, \rho)$
GossipFF	14%	14%	**43%**
GossipPB	14%	**21%**	0%
GossipPE	14%	**21%**	0%

Table III
GAIN IN TERMS OF EFFECTUAL FANOUT TO REACH $R = 99\%$

	$\mathcal{B}(N, p_N)$	$\mathcal{S}(N, m)$	$\mathcal{G}(N, \rho)$
GossipFF	40%	23%	**52%**
GossipPB	40%	**34%**	27%
GossipPE	40%	**34%**	31%

Table IV
GAIN IN TERMS OF EFFECTUAL FANOUT TO REACH $R = 80\%$

VII. RELATED WORK

In the previous sections, we have presented a thorough comparative study of the performance of three probabilistic gossip algorithms over three widespread topologies, thanks to the *effectual fanout*. There are several analysis and implementations in the literature but for a specific algorithm over one or two graphs or some given metrics. In the following, we provide a brief presentation of these works which are summarized in Table V.

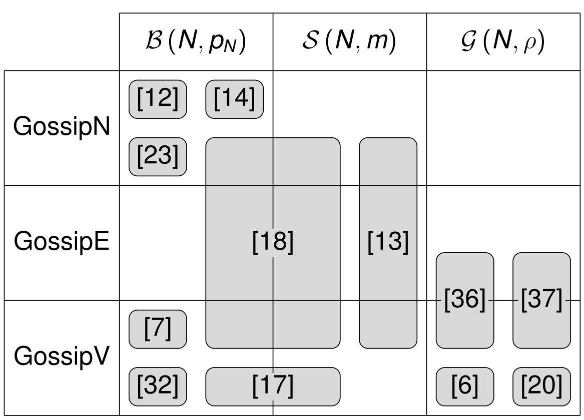

Table V
PREVIOUS STUDIES OF RANDOM GOSSIP ALGORITHMS

The reliability of the information dissemination is studied in [19] by applying *GossipFF* over $\mathcal{B}(N, p_N)$. it is assumed that the **fanout** of every site is always smaller than the number of its neighbors. The article mainly concludes that in a system with N sites to have the reliability equal to R, it requires to fix $fanout = -\ln\left(\frac{-\ln(R)}{N}\right)$. Results for *GossipFF* over $\mathcal{B}(N, p_N)$ which are based on simulations are also discussed in [10], [12]. However, the other gossip algorithms are not studied or compared in the articles. One the other hand, since the fanout is not linear to message complexity while the other two gossip algorithms take the probability as input parameter, the comparisons amongst them in function of their probabilistic input become difficult due to the lack of one generic parameter like effectual fanout.

The performance of *GossipPB* over $\mathcal{G}(N, \rho)$ is discussed and implemented in [4], [17] and it is theoretically analyzed over $\mathcal{B}(N, p_N)$ in [6]. The former is also studied in [31], aiming at answering how to choose p_v in order to reach high reliability. Besides the discussion about the reliability by percolation property over $\mathcal{B}(N, p_N)$, the asymptotic expressions in [24] with respect to the average number of messages and the average time required to complete network coverage are derived as well, showing the benefits of the properly choice of p_v. However, compared to our work, their efforts are focused on the performance of one gossip algorithm over a certain random topology, which can be considered as one aspect of our discussion. Over the two random topologies $\mathcal{S}(N, m)$ and $\mathcal{B}(N, p_N)$, the latency of a modified version of *GossipPB* algorithm which lets every site send message to one neighbor with certain probability several times is theoretically studied by SIS (Susceptible-Infective-Susceptible) model in [14].

According to the heuristic results firstly shown in [15], the performance of the three probabilistic gossip algorithms over $\mathcal{S}(N, m)$ is better than $\mathcal{B}(N, p_N)$. However, without effectual fanout, they cannot obtain the quantitative gains for all gossip algorithms in terms of message complexity to reach the same reliability.

GossipPE presents better performance than *GossipPB* Raman2009, which is studied in function of the system size (or the site degree) in [29]. Moreover, the three algorithms over $\mathcal{S}(N, m)$ are compared in the same way in [11]. In [28], the choice between *GossipPE* and *GossipPB* depends on the different application constraints over $\mathcal{G}(N, \rho)$. Compared to their work, we exploit the generic parameter, *effectual fanout*, which is linear to message complexity, and thus, the difference of all metrics can be fairly compared.

VIII. CONCLUSION

By exploiting the effectual fanout parameter, it is possible to finely observe the trade-off amongst dissemination reliability, message complexity, and latency for the three families of gossip algorithms with various kinds of probabilistic input over three graphs that characterize different random topologies. Thanks to this generic parameter which characterizes the sites mean dissemination, we have compared their performance in a quantitative way in this article.

On top of the simulation, we show that unlike the $fanout$

in *GossipFF*, the effectual fanout takes into account the degree distribution of the topology, which makes measures and performance evaluation more realistic and comparable over all the topologies and all the algorithms. Therefore, effectual fanout is a useful measurement parameter to uniform random networks with different degree distributions.

We have shown that in terms of reliability, *GossipFF* is the best algorithm on $\mathcal{G}(N, \rho)$ but the worst on $\mathcal{S}(N, m)$. All the algorithms have the same performance on $\mathcal{B}(N, p_N)$.

Results obtained in this study help in the decision of the most suitable combination between gossip algorithm and random topology to satisfy certain application requirements. It is also worth pointing out that the effectual fanout may theoretically analyze the performance, for instance, due to the equivalence of the three gossip algorithms over $\mathcal{B}(N, p_N)$, the relation between the effectual fanout and reliability can be studied in the similar way as in [19].

Our study presents the general choice for the combination between a gossip algorithm and a modeled random topology, while in fact, some results are already known in the literature for certain real applications: *GossipPE* and *GossipPB* are much more suitable for dissemination in TDMA-based networks; *GossipFF* proves its advantages relative to collisions and contentions in CDMA-based networks [22]; *GossipPE* is suitable for low-duty cycle sensor networks where few neighbors wake up simultaneously and also exhibits better performance than *GossipPB* in mobile ad hoc network with directional antennas [29].

We intend in a future work to extend our study to other gossip algorithms over various kinds of random topologies. Novel algorithm designs over other random topologies can also be evaluated and compared by this generic parameter. In other words, the effectual fanout proves to be a useful bridge amongst message redundancy, reliability, and latency for evaluating and studying novel algorithms.

References

[1] Omnet++: Discrete event simulation system *http://www.omnetpp.org/*.

[2] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, Jan. 2002.

[3] C. Avin. Distance graphs: From random geometric graphs to bernoulli graphs and between. *DIALM-POMC'08*, pages 71–78, 2008.

[4] B. Blywis, M. Günes, F. Juraschek, and S. Hofmann. Gossip routing in wireless mesh networks. *International Symposium on Personal Indoor and Mobile Radio Communications*, pages 1572–1577, 2010.

[5] B. Bollobás and O. Riordan. The diameter of a scale-free random graph. *Combinatorica*, 24:5–34, January 2004.

[6] S. Crisostomo, U. Schilcher, C. Bettstetter, and J. Barros. Analysis of probabilistic flooding: How do we choose the right coin? *International Conference on Communications*, pages 1–6, June 2009.

[7] N. Eggemann and S. Noble. The clustering coefficient of a scale-free random graph. *Discrete Applied Mathematics*, 159(10):953–965, 2011.

[8] P. Erdős and A. Renyi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5(17):17–60, 1960.

[9] P. Erdős and A. Rényi. On random graphs i. *Publicationes Mathematicae*, 6:290–297, 1959.

[10] P. T. Eugster, R. Guerraoui, S. B. Handurukande, A.-M. Kermarrec, and P. Kouznetsov. Lightweight probabilistic broadcast. *In ACM Transaction on Computer Systems*, 21:341–374, 2003.

[11] S. Ferretti and G. D'Angelo. Multiplayer online games over scale-free networks: a viable solution? *Proceedings of the 3rd International ICST Conference on Simulation Tools and Techniques*, 2010.

[12] D. Frey, R. Guerraouia, A.-M. Kermarrec, B. Koldehofe, M. Mogensen, M. Monod, and V. Quéma. Heterogeneous gossip. *International Conference on Middleware*, pages 42–61, 2009.

[13] R. Friedman and A. C. Viana. Gossiping on manets: The beauty and the beast. *Operating Systems Review*, 41(5):67–74, October 2007.

[14] A. Ganesh, L. Massoulié, and D. Towsley. The effect of network topology on the spread of epidemics. *INFOCOM'05*, pages 1455–1466, 2005.

[15] B. Garbinato, D. Rochat, and M. Tomassini. Impact of scale-free topologies on gossiping in ad hoc networks. *IEEE NCA*, pages 269–272, 2007.

[16] G. Grimmett. *Percolation*. Springer, 1989.

[17] Z. Haas, J. Halpern, and L. Li. Gossip-based ad hoc routing. *INFOCOM'02*, 3:1707–1716, 2002.

[18] J. Jetcheva, Y. Hu, D. Maltz, and D. Johnson. A simple protocol for multicast and broadcast in mobile ad hoc networks. *Internet Draft: draft-ietf-manet-simple-mbcast-01.txt*, 2001.

[19] A.-M. Kermarrec, L. Massoulie, and A. Ganesh. Probabilistic reliable dissemination in large-scale systems. *IEEE TPDS*, 3:248–258, March 2003.

[20] Z. Kong and E. M. Yeh. Characterization of the critical density for percolation in random geometric graphs. *IEEE ISIT*, 2007.

[21] J. Leitao, J. Pereira, and L. Rodrigues. Epidemic broadcast trees. *Int. Symposium on Reliable Distributed Systems*, pages 301–310, 2007.

[22] A. Muqattash and M. Krunz. Cdma-based mac protocol for wireless ad hoc networks. *Proceedings of MobiHoc*, pages 153–164, 2003.

[23] S.-Y. Ni, Y.-C. Tseng, Y.-S.Chen, and J.-P. Sheu. The broadcast storm problem in a mobile ad hoc network. *International Conference on Mobile Computing and Networking*, 8:151–162, 1999.

[24] K. Oikonomou, D. Kogias, and I. Stavrakakis. Probabilistic flooding for efficient information dissemination in random graph topologies. *Computer Networks*, 54:1615–1629, July 2010.

[25] M. Penrose. *Random Geometric Graphs*, volume 5. Oxford University Press, oxford studies in probability edition, May 2003.

[26] T. Philips, S. Panwar, and A. Tantawi. Connectivity properties of a packet radio network model. *IEEE Transactions on Information Theory*, pages 1044–1047, 1989.

[27] A. Polynikis. Random walks and scale-free networks. Master's thesis, University of York, UK, 2006.

[28] V. Raman and I. Gupta. Performance tradeoffs among percolation-based broadcast protocols in wireless sensor networks. *Workshop on Wireless Ad hoc and Sensor Networking*, pages 158–165, 2009.

[29] C.-C. Shen, Z. Huang, and C. Jaikaeo. Directional broadcast for mobile ad hoc networks with percolation theory. *IEEE Transactions on Mobile Computing*, 5(4):317–332, Apr. 2006.

[30] D. Ustebay, R. Castro, and M. Rabbat. Efficient decentralized approximation via selective gossip. *Journal of Selected Topics in Signal Processing*, PP(99):1, 2011.

[31] G. Vakulya. Energy efficient percolation-driven flood routing for large-scale sensor network. *IMCSIT*, 3:877–883, 2008.

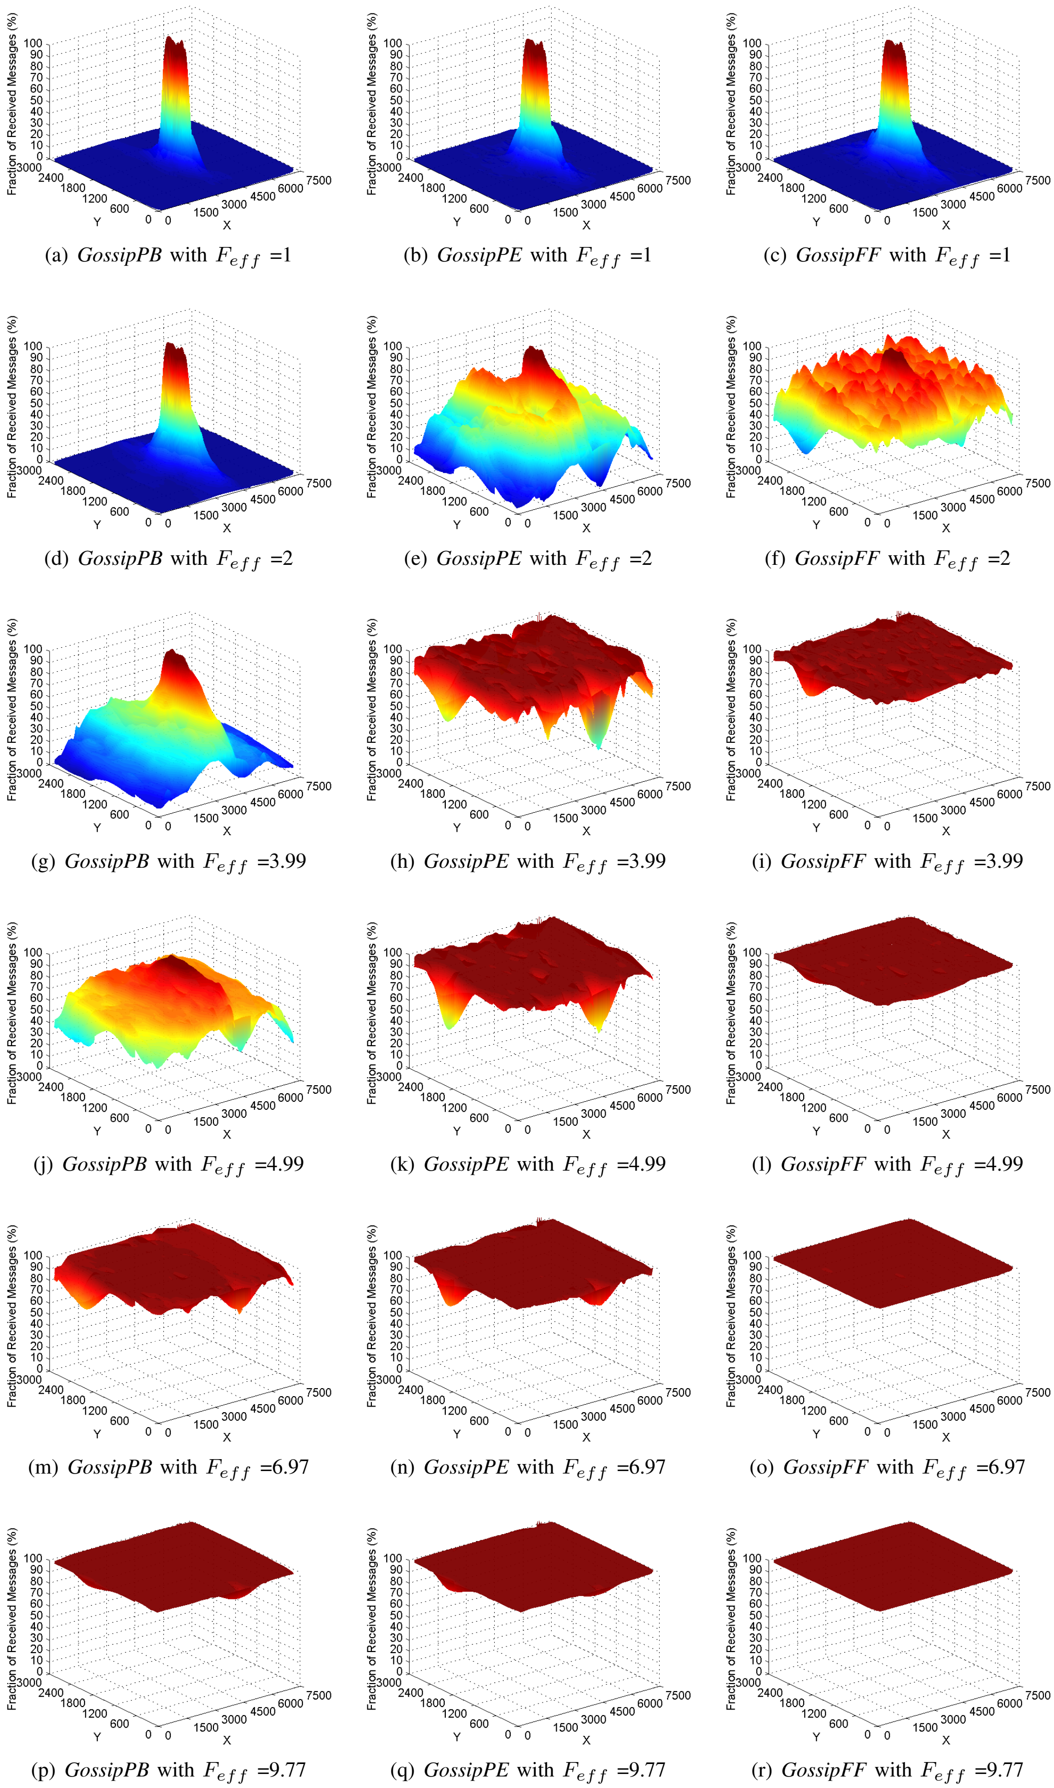

(a) *GossipPB* with F_{eff} =1
(b) *GossipPE* with F_{eff} =1
(c) *GossipFF* with F_{eff} =1
(d) *GossipPB* with F_{eff} =2
(e) *GossipPE* with F_{eff} =2
(f) *GossipFF* with F_{eff} =2
(g) *GossipPB* with F_{eff} =3.99
(h) *GossipPE* with F_{eff} =3.99
(i) *GossipFF* with F_{eff} =3.99
(j) *GossipPB* with F_{eff} =4.99
(k) *GossipPE* with F_{eff} =4.99
(l) *GossipFF* with F_{eff} =4.99
(m) *GossipPB* with F_{eff} =6.97
(n) *GossipPE* with F_{eff} =6.97
(o) *GossipFF* with F_{eff} =6.97
(p) *GossipPB* with F_{eff} =9.77
(q) *GossipPE* with F_{eff} =9.77
(r) *GossipFF* with F_{eff} =9.77

Figure 6. The message reception of every site.

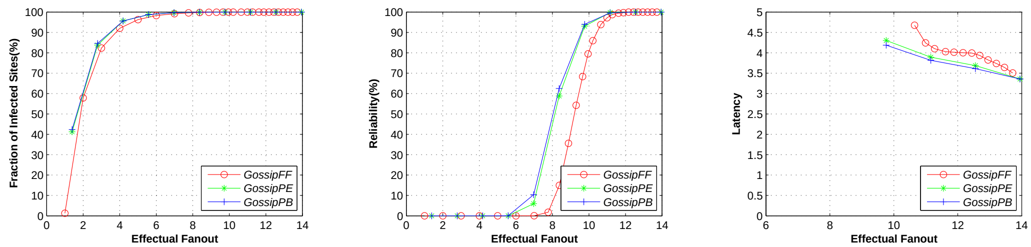

(a) Infected Sites Over $\mathcal{S}(N, m)$ (b) Reliability Over $\mathcal{S}(N, m)$ (c) Latency Over $\mathcal{S}(N, m)$

Figure 7. Performance comparison of algorithms over $\mathcal{S}(N, m)$.

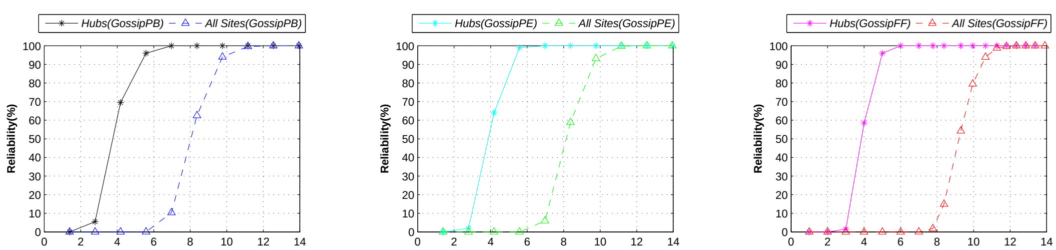

(a) Reliability by *GossipPB* (b) Reliability by *GossipPE* (c) Reliability by *GossipFF*

Figure 8. The reliability of all hubs and all sites by gossip algorithms over $\mathcal{S}(N, m)$