

On Modeling Fault Tolerance of Gossip-Based Reliable Multicast Protocols

Xiaopeng Fan¹, Jiannong Cao¹, Weigang Wu¹, Michel Raynal²,

1 Department of Computing, Hong Kong Polytechnic University, Kowloon, Hong Kong

2 IRISA, Campus de Beaulieu, 35042 Rennes Cedex, France

{csxpfan, csjcao, cswgwu}@comp.polyu.edu.hk, raynal@irisa.fr

Abstract

Gossiping has been widely used for disseminating data in large scale networks. Existing works have mainly focused on the design of gossip-based protocols but few have been reported on developing models for analyzing the fault tolerance property of these protocols. In this paper, we propose a general gossiping algorithm and develop a mathematical model based on generalized random graphs for evaluating the reliability of gossiping, i.e., to what extent gossip-based protocols can tolerate node failures, yet guarantee the specified message delivery. We analytically derive the maximum ratio of failed nodes that can be tolerated without reducing the required degree of reliability. We also investigate the impact of the parameters, namely the fanout distribution and the nonfailed member ratio, on the protocol reliability. Simulations have been carried out to validate the effectiveness of our analytic model in terms of the reliability of gossiping and the success of gossiping. The results obtained can be used to guide the design of fault tolerant gossip-based protocols.

1. Introduction

Reliable group communication protocols are essential for designing distributed systems and applications such as publish/subscribe systems [1], distributed databases [2], consistency management [3], and distributed failure detection [4]. In recent years reliable multicast protocols for the large-scale networks have been the focus of research in group communication. Traditional solutions applicable in small-scale settings are not scalable and reliable in large distributed systems. How to design multicast protocols guaranteeing specified reliability in large-scale systems has become a challenging problem for researchers.

Existing multicast protocols guarantee one of the three types of reliability: strong reliability, best-effort

reliability and probabilistic reliability [5]. Compared with strong reliability and best-effort reliability, probabilistic reliability does not always guarantee atomicity but can provide message delivery guarantee with some required probability. For example, Bimodal Multicast [5] provides a bimodal delivery guarantee which changes the traditional “all or nothing” guarantee to the “almost all or almost none” guarantee.

Gossiping is one of the most important techniques to provide probabilistic reliability in reliable multicast. Gossip-based multicast protocols rely on a peer-to-peer interaction model for multicasting a message, and they are scalable since the load is distributed among all participating nodes. Redundant messages are used to achieve reliability and fault tolerance. A few pioneering works on gossiping have been done for both wired and wireless networks. In wired networks, work can be found on data dissemination [6], consistency management in replicated databases [2], and failure detection [4]. In wireless networks, gossip-based protocols have been proposed for multicast in mobile ad hoc networks (MANETs). A seminal approach is the Anonymous Gossip (AG) protocol [7], which is a descendant of the pbcast [5] protocol. The Route Driven Gossip (RDG) [8] protocol uses a pure gossip scheme, by which messages, negative acknowledgments, and membership information are gossiped uniformly without requiring an underlying multicast primitive.

In this paper, we are interested in analyzing the impact of the fanout distribution and node failures on the performance of gossiping protocols. Existing works have mainly focused on the design of gossip-based protocols, but not many works can be found on developing mathematic models for analyzing the properties of the protocols. In some existing works, the process of gossiping is modeled as a Markov Chain, as proposed in pbcast [5] and RDG [8]. These models need to be simplified due to their intractability, but the simplification will affect the accuracy of the reliability analysis. Actually, only upper bounds or lower bounds on the reliability can be obtained. Another modeling

technique is originated from the observation on the similarity between gossiping and the epidemic spreading of disease. Authors of [9] propose an epidemic model for analyzing the Local Retransmission-based Gossip (LRG) protocol. However, the model cannot explain how to obtain the optimal value of the probability with which a node gossips to other nodes, and does not consider node failures. Recently, random graphs [10] have been used as a promising technique to model gossiping. The work reported in [6] by Microsoft proposes to model the success of gossiping as the existence of a directed path from the source node to every other node in a random graph. However, the model only consider the success of gossiping, i.e., all of the members receive the message, but not the probability that one node receives the message during gossiping. We think the latter is also important for analyzing the performance of gossiping and consider both of cases in our model.

In this paper, we first describe a general gossiping algorithm which differs from existing algorithm, allows each node to generate a random number of gossiping targets by following a specified probability distribution. In traditional gossiping algorithms, each node normally has a fixed number of gossiping targets. Targeted at this proposed algorithm, we develop a mathematical model to analyze its fault tolerance property using the generalized random graph theory [11]. Observing that the process of generating a random graph is similar to the process of gossiping a message in a multicast group, we use the size of the giant component in a random graph to represent the probabilistic reliability of gossiping, in the sense that nodes in the giant component can be reached by the source node with a high probability. Taking node failures into consideration, we analyze the performance of gossiping in terms of the reliability and the success of gossiping.

Our generalized random graph model has two advantages. First, using generalized random graphs enables us to study various fanout distributions, including the Poisson fanout distribution, which has been adopted in most of the previous works. We match our mathematic model to our proposed algorithm harmoniously. The varieties of gossiping are also enriched by the algorithm. Second, our mathematical model allows us to derive the critical point of gossiping at which the maximum ratio of failed nodes can be tolerated yet the required reliability can still be guaranteed. Analytically, we study the impact of the fanout distribution and the nonfailed member ratio on the performance of gossiping. We have validated the analytical results by using simulations.

The rest of the paper is organized as follows: Section 2 discusses the related work. Section 3

provides the network model, the general gossiping algorithm, and the preliminaries for the theory of generalized random graphs. In Section 4, we present a mathematical model for analyzing fault tolerance of our gossiping algorithm. Simulation results are reported in Section 5. Finally, Section 6 concludes this paper.

2. Related Work

In this section, we briefly review the previous work on developing mathematic models of gossiping. Three different modeling approaches have been used: the recurrence model, the epidemic model, and the random graph model.

In pbcast [5], the analysis shows how to calculate the bimodal delivery distribution for a given networking setting. The authors derive a recurrence relationship between successive gossiping rounds of the protocol. However, due to its complexity, the model cannot give the accurate value of the reliability of gossiping. Analysis based on this model only calculates an upper bound on the probability in round t that s_{t+1} nodes will receive the gossip message in the next round $t+1$. It does not show how to find a proper number of rounds required in gossiping.

The second approach is based on the epidemic model [9]. Two mechanisms, Local Retransmission and Gossiping (LRG), are combined to provide the high reliability of data delivery. Since the gossiping process is similar to the spreading of epidemic diseases, the authors use the so-called SI model in epidemiology to analyze LRG. In this model, the balance equations are developed to describe the process of spreading messages among GCHs (Group Cluster Heads). However, the model did not take message losses and node failures into consideration.

More recently, the random graph theory [10] has been used to model gossiping. The seminal work is the model by Microsoft [6], which aims at establishing the relationship between the success of the gossiping protocols and the key gossips parameters, including fanout and failure rate. The model considers the presence of arc $\{x, y\}$ in a random graph $\zeta(n, p_n)$ as saying that x gossips message to y . The success of gossiping means the existence of a directed path from the source node s to every other node in the random graph. The probability of the success of gossiping is denoted by $\pi(p_n, n)$. It is proved that the limit of

$\pi(p_n, n)$ is $e^{-e^{-c}}$ if p_n is equal to $[\log n + c + o(1)]/n$, where c is a constant. If the proportion of the failed nodes is ε , that is, $n' = (1 - \varepsilon)n$, gossiping succeeds

with the probability $e^{-e^{-c}}$ if $p_n = [\log n' + c + o(1)]/n'$. Although the success of gossiping, i.e., all of the group members receive the message, is important, we still need to know the probability that one node receives the message during gossiping if we cannot guarantee such a strong status as the success of gossiping practically. Both of them are discussed in detail by our mathematic model.

Compared with Microsoft's work [6], our model proposed in this paper have three advantages. First, it proposes a novel way to represent the reliability of gossiping by the size of the giant component [11] in a generalized random graph, which provides a simpler way to describe the reliability of gossiping. Second, we discuss not only the success of gossiping, but also the reliability of gossiping, i.e., the percentage of nonfailed nodes that receive the message. Third, analysis using our model can be performed for various fanout distributions, rather than only the Poisson distribution. Gossiping tailored for different applications over various types of overlays or physical topologies may need to use different fanout distributions in order to improve the performance of gossip-based protocols.

3. Preliminaries

In this section, we first introduce our system model and describe a general gossiping algorithm. Then we provide the preliminaries on the theory of generalized random graphs.

In our system model, a multicast group G is composed of n members, which have the interest to share the same message m . Each member has a unique ID. We consider a fail-stop failure model where failed members will not gossip messages they receive and they fail only by crashes. Moreover, we assume the source node never fails. In real applications, we assume that a scalable membership protocol is available, such as [12], [13], which can be applied to gossip-based reliable multicast protocols in large-scale systems. Membership is beyond the scope of this paper and will not be discussed further.

We use a general gossiping algorithm as shown in Figure 1, which allows for various distributions of the fanout of the nodes. When a member receives the message m for the first time, it generates a random number f_i by following a specified probability distribution P . Then the node chooses f_i gossip targets from its own membership view and sends the message out. If a member receives the message again, it discards it immediately.

The reliability of gossiping is defined as the ratio of the number of nonfailed members that received the message m to the total number of nonfailed members

in the group G . We denote the reliability of gossiping as $R(q, P)$, the probability that a nonfailed member can receive the message m after one execution of our algorithm. The success of gossiping is defined as all of the nonfailed members received the message m at least once after t executions of our gossiping algorithm, denoted by $S(q, P, t)$. In this paper, we focus on the relationship between the parameters of the gossiping algorithm and the reliability of gossip-based multicast protocols. The key parameters in gossiping are listed as follows:

- P : *Fanout Distribution*, the probability distribution of the fanout of members
- q : *Nonfailed Member Ratio*, the ratio of the number of the nonfailed members to the number of the total members.

Algorithm for each node in multicast group G

Upon member i receiving the message m for the first time

{

Member i generates a random number f_i by following a specified probability distribution P

Member i selects f_i nodes uniformly at random from its membership view

Member i sends the message m to the selected f_i nodes

}

Fig. 1 **The gossiping algorithm**

Compared with the traditional Poisson random graph model, generalized random graph [11] [14], which has been previously applied in Physics, is a more general model for random graphs. It is applicable to arbitrary degree distribution in a random graph. Before we introduce our mathematical model for gossiping, we briefly introduce some fundamental concepts of generalized random graphs: *Degree Distribution*, *Component*, *Phase Transition*, and *Giant Component*.

Degree Distribution denotes the probability distribution of the degrees of nodes in a generalized random graph. A *Component* is a set of nodes that can reach each other along the paths on the graph. A *Phase Transition* refers to the phenomenon that, while a random graph with n nodes and a certain number of edges is unlikely to have one special property, a random graph just with a few more edges is very likely to have this property. A good example of a phase transition is of the critical point of the connectivity of a random graph. The critical point means the point at which the connectivity of a random graph grows dramatically. The *Giant Component* is the biggest component formed after a phase transition happens. It

has a size of order at least $n^{2/3}$, while the sizes of other components are of order at most $n^{2/3}/2$.

In our gossip algorithm, the number of each member's gossip targets is a random variable and these random variables are independent and identically-distributed random variables. In fact, it is similar to the case that we draw samples from the total population but without replacing any of them, because each nonfailed member only gossips once. The execution of the gossiping algorithm dies out when the nonfailed members that received the message m are in the same connected component. Therefore, the distribution of the fanout has the most important impact on the reliability of gossiping and the success of gossiping.

4. A Fault-Tolerant Gossip Model

In this section, we first propose an analytical model for fault tolerant gossiping by using the theory of generalized random graphs. Here the term of "fault-tolerant" means that if a proposed gossiping algorithm aims to achieve a required reliability, it should take node failures into consideration. Then we show how to use this model to analyze the performance of the gossiping algorithm by taking the Poisson fanout distribution as an example.

4.1 Model Definition

A gossiping model $Gossip(n, P, q)$ consists of n members to participate in gossiping with the fanout distribution P . Only a ratio q of all of the members can work correctly and other nodes may fail by crashes during gossiping. We consider two cases of failures which are treated the same. Members may crash either before receiving the message or after receiving the message but not yet forwarding it to others. P is the fanout distribution of nonfailed members that participate in gossiping. As mentioned before, we assume that the source node that initiates gossiping never fails. We use the terms "source member" and "source node" interchangeably.

Let $\zeta(n, P)$ be the space of generalized random graphs generated by gossiping, containing n nodes in the group G , and each node chooses its gossiping targets from its own membership view. Let p_k be the probability that a randomly chosen node from one element in $\zeta(n, P)$ has the degree k , and q_k be the probability that a node with the degree k is also a nonfailed node. The degree distribution in the random graph $\zeta(n, P)$ can be generated by the following generating function [15]:

$$F_0(x) = \sum_{k=0}^{\infty} p_k q_k x^k \quad (1)$$

In the above model, p_k can be any probability distribution. But we investigate the special case of uniform probability for q_k because we assume the setting for each member gossiping is the same one. We set $q_k = q$ for all k , i.e. all of the nodes fail with the same probability $(1-q)$. It is trivial that the number of nonfailed nodes equals to n^*q .

4.2 Analysis of Gossiping

The methodology is to investigate the properties of generalized random graphs to analyze the performance of the gossiping algorithm. We consider two of the most important problems: how to evaluate the reliability of gossiping and how to guarantee the success of gossiping. There are n members in the group G and the number of the nonfailed members is denoted by $n_{nonfailed} = [n^*q]$. We define n_{rece} as the number of nonfailed members that receive the message m after one execution of the algorithm. The reliability of gossiping $R(q, P)$ can be defined as follows:

$$R(q, P) = n_{rece} / n_{nonfailed}$$

The number of nonfailed members that receive the message m at least one time after t executions of the algorithm is referred to as n_{rece}^t . The success of gossiping $S(q, P, t)$ can be defined as follows:

$$Pr(S(q, P, t)) = Pr(n_{rece}^t = n_{nonfailed}),$$

where all of the nonfailed members receive the message m at least one time after t executions of the algorithm.

(1) Reliability of Gossiping: $R(q, P)$

Since the giant component changes the connectivity of random graph, the probability that a randomly chosen node belongs to this component is increased dramatically. With the size of the giant component growing, the probability that nodes receive the message is increased, which means more and more members can receive the message m sent from the source node in gossiping.

Firstly, the condition for the appearance of the giant component can be obtained in the following steps. According to the generalized random graph theory, the mean size $\langle s \rangle$ of the components in the random graph $\zeta(n, P)$ is

$$\langle s \rangle = q \left[1 + \frac{qG_0'(1)}{1 - qG_1'(1)} \right] \quad (2)$$

where $G_0(x) = \sum_{k=0}^{\infty} p_k x^k$ is defined as the generating function for the probability distribution of nodes

degree k , and $G_1(x) = \frac{G'_0(x)}{G'_0(1)}$ is the generating function

of the probability distribution of number of outgoing edges. Eq. (2) diverges where the equation $1 - qG'_1(1) = 0$ is satisfied, which is also the critical point at which a random graph achieves the giant component. According to the result above, the nonfailed member ratio at the critical point is

$$q_c = \frac{1}{G'_1(1)} \quad (3)$$

q_c is the critical point at which the random graph begins to improve its connectivity dramatically.

Secondly, we refer to S as the size of the giant component, which means the ratio of the number of the nonfailed nodes in the giant component to the total number of nonfailed nodes in the random graph. S can be calculated by the following equation:

$$S = F_0(1) - F_0(u) \quad (4)$$

where u is the solution of the self-consistency condition $u = 1 - F_1(1) - F_1(u)$, and $F_1(x)$ is defined as $F_0(x) / G'_0(1)$ [15].

(2) Probability of Success of Gossiping: $Pr(S(q, P, t))$

The success of gossiping is the result that all of the nonfailed group members receive the message m . In one execution of our gossiping algorithm, we can increase the probability of gossiping success by increasing the fanout of members. But this method is not a pragmatic one in the implementation of real applications. For example, one node only has one physical neighbor but it still needs to send the message m to all of the other nodes at one execution. Therefore, we consider another means to increase the probability of the success of gossiping by increasing the number of executions of our gossiping algorithm, in order to guarantee the required probability of the success of gossiping.

In the repeated executions, each execution can be viewed as one independent Bernoulli trial. So t times of executions can be considered as a t times Bernoulli trials. We define X as the number of executions in which a nonfailed member receives the message m during t executions. We do not consider how many times for each nonfailed member to receive the message m in one execution. We use p_r to denote the requirement of the reliability $R(q, P)$, and it is obvious that X follows a Binomial distribution $B(t, p_r)$. The distribution of X is in the following:

$$P(X = k) = C_t^k (p_r)^k (1 - p_r)^{t-k}, \quad k = 0, 1, 2, \dots, t$$

The probability of the success of gossiping $S(q, P, t)$ can be calculated by the following:

$$\Pr(S(q, P, t)) = P(X \geq 1) = 1 - (1 - p_r)^t \quad (5)$$

If the requirement for the success of gossiping is denoted by the probability p_s , we can obtain the requirement of t as follows:

$$t \geq \lg(1 - p_s) / \lg(1 - p_r) \quad t \in N \quad (6)$$

4.3 Case Study: Poisson Fanout Distribution

In this section, we take the Poisson distribution as an example of fanout distribution to show how to apply our mathematical model in the performance analysis of gossiping.

The fanout distribution is specified by a Poisson distribution $Po(z)$, where z is the mean of Poisson distribution $Po(z)$, and is also the average fanout. Then, the gossiping model can be defined as $Gossip(n, Po(z), q)$. This gossiping model can be modelled by a random graph model $\zeta(n, Po(z))$. Let p_k be the probability that a randomly chosen node from $\zeta(n, Po(z))$ has the degree k , and q be the nonfailed node ratio. It is important to notice that, although the distribution of node degree may be changed by node failures, it is always a Poisson distribution but with a smaller mean fanout q^*z [15].

We can obtain the following generating functions for $\zeta(n, Po(z))$:

$$F_0(x) = \sum_{k=0}^{\infty} p_k q x^k = e^{zq(x-1)} \quad (7)$$

$$G_0(x) = \sum_{k=0}^{\infty} p_k x^k = e^{z(x-1)} \quad (8)$$

$$G_1(x) = \frac{G'_0(x)}{G'_0(1)} = e^{z(x-1)} \quad (9)$$

According to Eq. 3, the critical point q_c can be obtained by

$$q_c = \frac{1}{G'_1(1)} = \frac{1}{z}$$

This means that, to the guarantee on the reliability of gossiping, the nonfailed member ratio q should be greater than $1/z$, i.e.:

$$q > 1/z \quad (10)$$

It is trivial that $G_0(x) = G_1(x) = e^{z(x-1)}$ in case of the Poisson distribution. Following Eq. (4), we can obtain the size of the giant component by

$$S = 1 - e^{-zqS} \quad (11)$$

Eq. (11) shows that the reliability of gossiping $R(q, Po(z))$ can be improved if we increase the fanout z or q .

Then, given the reliability of gossiping (represented by S) and the nonfailed node ratio q , the mean fanout z of the Poisson distribution can be obtained as follows:

$$z = \ln(1 - S) / (-qS) \quad (12)$$

Fig. 2 shows the numerical results of z against S under various q . With these results, we can determine the proper mean fanout for the Poisson distribution. But remember that Eq. (10) should still be held. The reliability of gossiping ranges from 0.1111 to 0.9999.




Fig. 2 Mean fanout vs. Reliability of Gossiping under various nonfailed node ratio




Fig. 3 Minimum times of executions for the required probability of gossiping success

Since the reliability of gossiping can be evaluated by the size of the giant component S , the condition for the success of gossiping $S(q, Po(z), t)$ in Eq. (6) can be revised as follows:

$$t \geq \lg(1 - p_s) / \lg(1 - S) \quad t \in N.$$

Fig. 3 depicts the analytical results of the minimum number of executions with a specified probability of the success of gossiping.

5. Simulations

To examine the effectiveness of our analytic model, we have carried out extensive simulations. We evaluate the performance of our gossiping algorithm according to the following metrics:

- The reliability of gossiping
- The success of gossiping

We use MATLAB 7.0 to implement and execute our gossiping algorithm. We evaluate the performance of the algorithm with two different group sizes of the groups 1000 and 5000 members. The key parameters, i.e., the fanout distribution and the nonfailed member ratio, are varied to evaluate their impact. In our simulation, we take Poisson distribution as an example for both the simulations and our analysis. Compared with the analytical results obtained by our mathematical model, the simulation results are well consistent.

5.1 Reliability of Gossiping




Fig. 4a Reliability in a 1000 nodes group

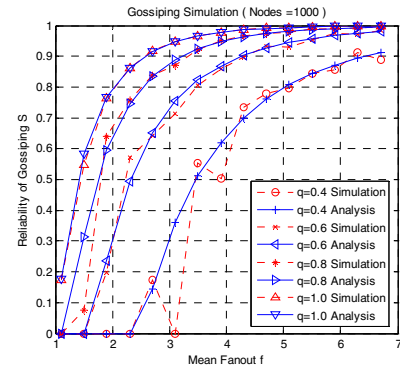


Fig. 4b Reliability in a 1000 nodes group

We set the nonfailed member ratio q to be 0.1, 0.2, 0.3, ..., 1.0 respectively. The mean fanout f for the fanout distribution is varied from 1.10 to 6.7 with an incremental step 0.4. The reason why we select this range is the value of the reliability is almost covered from 0 to 1. For each pair of $\{f, q\}$, we run our gossiping algorithm 20 times and report the average results of the reliability of gossiping. In addition, we calculate the size of giant component for each case.

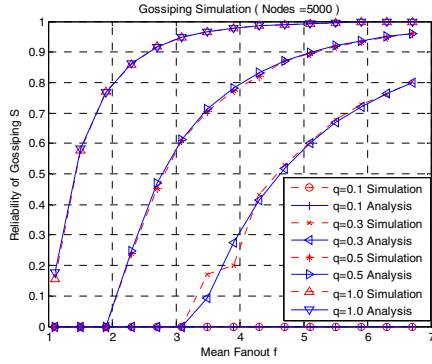


Fig. 5a Reliability in a 5000 nodes group

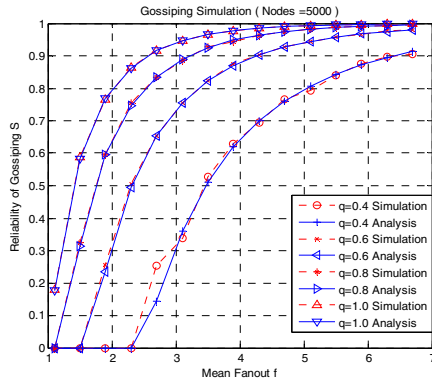


Fig. 5b Reliability in a 5000 nodes group

Figs. 4 and 5 depict the results in simulation and analysis for configurations with 1000 and 5000 nodes respectively. To see the results clearly, we divide each figure into two plots by a group of different q . In each plot, each dotted line denotes the simulations, while the continuous line represents the size of giant components using our mathematical model solved by Eq. (11).

We first observe that all of the critical points for each fanout are held under the condition that the nonfailed member ratio q should be greater than the reciprocal of the mean fanout f . For each fanout in our simulation, the reliability of gossiping can be guaranteed under the above condition. Fig. 4 also shows that the results of simulations tally with the analytical results except very few points. The curves in Fig. 5 are very similar to those in Fig. 4. However, the simulation results tally with the analytical results better than in Fig. 4, which indicates that our modeling works better in larger scale systems.

5.2 Success of Gossiping

Besides reliability, we also measure the success of gossiping. We select two pairs of key parameters $\{f, q\}$ as follows: $\{4.0, 0.9\}$, and $\{6.0, 0.6\}$. The requirement for the success of gossiping is set to 0.999, the same value as in our analysis. For each pair of parameters,

we run our gossiping algorithm for 20 times in one simulation, and each simulation is repeated for 100 times. Then we report the distribution of the number X , i.e. the number of gossiping succeeds among 20 executions. If X is approximately follows a binomial distribution $B(20, R(q, Po(z)))$, this means the calculation in Eq. (6) is valid.

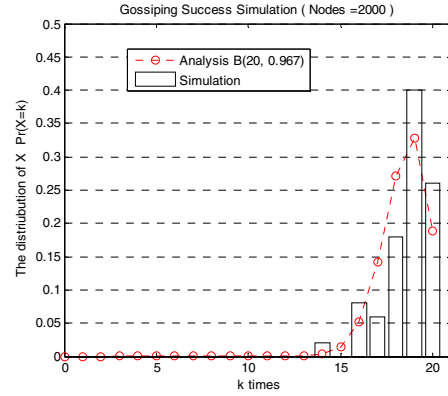


Fig. 6 The distribution of Gossiping Success with $f=4.0, q=0.9$

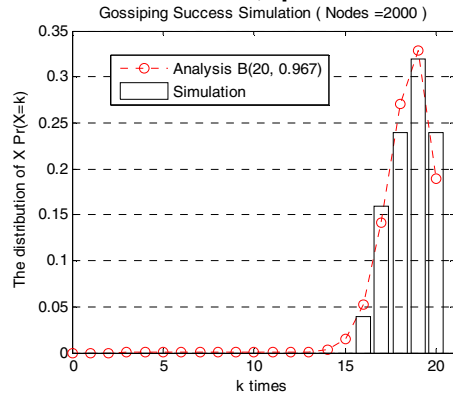


Fig. 7 The distribution of Gossiping Success with $f=6.0, q=0.6$

Figs. 6 and 7 plot the results of simulations and analysis in a group with 2000 nodes. In the figure, each bar represents the simulation result of the probability $Pr(X=k)$ where k ranges from 0 to 20, while the continuous line represents the value of $Pr(X=k)$ from $X \sim B(20, R(q, Po(z)))$. According to Eq. (6), we can obtain the required number of executions as follows:

$$t \geq \lg(1 - 0.999) / \lg(1 - 0.967) \quad t \in N$$

t should be greater than three. Figs. 6 and 7 show the simulation results tally with our analytic results well.

It is interesting to notice that the gossiping with $\{4.0, 0.9\}$ and $\{6.0, 0.6\}$ can obtain the same reliability of gossiping in one execution as 0.967 because the product of $f \cdot q$ are the same one. However, their corresponding distributions of gossiping success are not exactly identical. This is because the mean fanout

and the nonfailed node ratio have different impact factors on the probability of the success of gossiping.

6. Conclusions

Based on the generalized random graph theory, we develop a mathematical model analyzing the performance of the gossiping algorithm in terms of the reliability of gossiping, and the success of gossiping. We focus on the fault tolerance of gossiping by taking node failures into consideration. We propose to represent the reliability of gossiping by using the size of the giant component in a random graph for the first time. Our model can be resolved by the generalized random graph theory and derive the relationship between the parameters of gossiping, and the reliability of gossiping and the success of gossiping. Our model shows that there exists a threshold value of the number of the nonfailed nodes ratio for guaranteeing a specific reliability in gossiping. We have carried out extensive simulations to validate our proposed model. The simulation results tally with our analytic results very well. Therefore, our analytic model is effective and accurate.

7. Acknowledgements

This research is partially supported by the Hong Kong Polytechnic University under the CERG grant B-Q937(PolyU 5105/05) and French/Hong Kong Joint Research Scheme under grant F-HK16/05T.

8. References

- [1] P. T. Eugster, R. Guerraoui, S. B. Handurukande, P. Kouznetsov, and A. Kermarrec, "Lightweight probabilistic broadcast", *ACM Tran. Computer*, (4): pp. 341-374 (2003).
- [2] A.J. Demers, D.H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry, "Epidemic Algorithms for Replicated Database Maintenance," *Proc. Sixth Ann. ACM Symp. Principles of Distributed Computing (PODC)*, pp. 1-12, Aug. 1987.
- [3] R. Golding and K. Taylor, "Group Membership in the Epidemic Style", *Technical Report UCSC-CRL-92-13*, Dept. of Computer Science, Univ. of California, Santa Cruz, 1992.

- [4] R. Renesse, Y. Minsky, and M. Hayden, "A Gossip-Style Failure Detection Service", *Proceedings of IFIP International Conference on Distributed Systems Platforms and Open Distributed Processing*, Middleware'98.
- [5] Birman, Ken, Hayden, Mark, Ozkasap, Oznur, Xiao, Zhen, Budi, Mihai and Minsky, Yaron "Bimodal Multicast", *ACM Transactions on Computer Systems*, Vol. 17, No. 2, pp 41-88, May, 1999.
- [6] A.-M. Kermarrec, L. Massoulié', and A.J. Ganesh, "Probabilistic Reliable Dissemination in Large-Scale Systems", *IEEE Trans. Parallel and Distributed Systems*, VOL. 14, NO. 3, March, 2003.
- [7] R. Chandra, V. Ramasubramanian, and K. Birman, "Anonymous gossip: improving multicast reliability in mobile ad hoc networks," in Proc. 21st Int. Conf. Distributed Computing Systems (ICDCS), 2001, pp. 275-283.
- [8] J. Luo, P. T. Eugster, and J. P. Hubaux, "Route driven gossip: Probabilistic reliable multicast in ad hoc networks", *Proc. IEEE INFOCOM*, Apr. 2003, pp. 2229-2239.
- [9] W. Jia, D. Lu, G. Wang, L. Zhang and W. Lu, "Local Retransmission-based Gossip Protocol in Mobile Ad Hoc Networks", *IEEE WCNC 2007*, pp. 4244-4249.
- [10] B. Bollobás, *Modern Graph Theory*, Springer-Verlag, New York, pp. 215-251.
- [11] Newman, M. E. J., Strogatz, S. H., and Watts, D. J., "Random graphs with arbitrary degree distributions and their applications", *Phys. Rev. E* 64, 026118 (2001).
- [12] A.J. Ganesh, A.-M. Kermarrec, and L. Massoulié', "SCAMP: Peer-to-Peer Lightweight Membership Service for Large-Scale Group Communication", *Proc. Third Int'l Workshop Networked Group Comm. (NGC)*, Nov. 2001.
- [13] A. Allavena, A. Demers and J. Hopcroft, "Correctness of a Gossip-based Membership Protocol", *Proc. 24th ACM Symposium on the Principle of Distributed Computing (PODC)*, pp. 292-301, 2005.
- [14] M. E. J. Newman, "The Structure and Function of Complex Networks", *SIAM Review* 45, pp. 167-256, 2003.
- [15] D. Callaway, M. Newman, S. Strogatz and D. Watts. "Network Robustness and Fragility: Percolation on Random Graphs", *Phys. Rev. Lett.* 85, pp. 5468-5471. 2000