A Flooding Method for Exchanging Routing Information in Power–Law Networks

Nobutaka Makino†, Shin'ichi Arakawa‡, and Masayuki Murata†
†Graduate School of Information Science and Technology, Osaka University, 1-5 Yamadaoka, Suita, Osaka 560-0871, Japan
‡Graduate School of Economics, Osaka University, 1-7 Machikaneyama, Toyonaka, Osaka 560-0043, Japan
Email: {n-makino, murata}@ist.osaka-u.ac.jp, arakawa@econ.osaka-u.ac.jp

Abstract— **Recent measurement studies of the Internet topology show that the connectivities of nodes exhibit power–law attributes. This topology has two main characteristics: many nodes have a small number of links while a few nodes have a large number of links; and there are fewer hop–counts between nodes. Previous studies of routing mechanisms have evaluated their validity or effectiveness in relatively small networks. However, an evaluation of power–law networks is also needed to clarify their actual validity. In this paper, we evaluate some of the flooding mechanisms used in routing protocols. Our simulation results show that in a power–law network, the flooding mechanism does not scale well due to the concentration of message, whereas random networks that do not have power–law attributes actually scale well. To reduce the concentration of messages, we propose an efficient flooding method for power–law networks. Our method uses probabilistic flooding in which each node relays routing information with a certain probability. Routing information is also exchanged periodically to prevent information mismatches between nodes. The simulation results showed that, compared to conventional flooding approaches, our method reduced the amount of traffic by 50%.**

I. Introduction

Flooding is used to exchange routing information on the Internet. For example, in the OSPF (Open Shortest Path First) protocol [1], a node that acquires a change in link status distributes messages that include link–state information to its neighbor nodes. Each neighbor node that receives the link–state information redistributes the information to its corresponding neighbor nodes. The BGP (Border Gateway Protocol) also uses a flooding mechanism to exchange routing information. Each node establishes a TCP connection to each neighbor node and then transfers the routing table. In both protocols, since the distribution or exchange of routing information is based on a flooding mechanism, the amount of traffic involved in exchanging route information is becoming a critical problem as the number of nodes connected to the Internet increases [2].

Recent measurement studies on the topology of the Internet show that the connectivities of nodes exhibit power–law attributes [3]. That is, the probability $p(k)$ that a node is connected to k other nodes follows $p(k) \sim k^{-\gamma}$ (γ is a constant). In recent years, a considerable number of studies have investigated power–law networks whose degree distribution follows the power–law. Most research on power–law networks has focused on investigating how to model the topology of the Internet. A theoretical examination of the characteristics of the topology is also presented in [4]. The power–law network has two main characteristics: (1) many nodes have a small number of links while a few nodes have a large number of links, and (2) the number of hop–counts between nodes is reduced (*small–world* property) [4, 5].

The second characteristic promotes faster propagation of information, which is an advantage in exchanging route information. However, because of the first characteristic, if several nodes perform flooding at the same time or at almost the same time, control messages concentrate at the hub node. For example, if a node fails, flooding starts from all of the neighbor nodes, which causes sudden traffic congestion in the network. It is also likely that this tendency will increase as the number of nodes in the network increases because the number of links connected to the hub–node in turn increases. Previous evaluations of routing or flooding mechanisms have been performed on random networks. However, an evaluation of power–law networks is also needed to clarify their actual validity.

In this paper, we first describe an evaluation of the characteristic of the number of control messages generated by conventional flooding methods. We show that a conventional flooding method generates an enormous number of duplicate control messages at the hub–node. Based on this observation, we propose a new flooding method that generates fewer duplicate messages based on a consideration of the topological characteristics of power–law networks. With conventional flooding methods, a control message is duplicated at an intermediate node and delivered to its neighbor nodes. If a node has multiple routes to another node, the later node receives the same message propagated via different routes. Our method reduces the number of control messages that each node relays by using a probabilistic relay method. The method also ensures that information is delivered to all nodes by periodically exchanging information between adjacent nodes.

Section II of this paper outlines conventional flooding methods. Section III evaluates the number of control messages generated by a conventional flooding method over a power–law network. In Section IV, we present a new flooding method based on the topological characteristics of power–law networks, and we shows the effectiveness of the proposed method by comparing it with conventional flooding methods in Section V. Our conclusions are presented in Section VI.

II. Related Works

Various methods have been proposed to reduce the number of control messages caused by flooding [6-9] In Ref. [7], the FSLS (Fuzzy Sighted Link State) method was proposed. This method restricts the forwarding area of control messages by setting a TTL (Time–To–Live) value. The TTL is the maximum number of hops that a message can survive. A large TTL value means that the control message is delivered to a large number of nodes, while a small TTL value restricts its delivery to a limited number of nodes. The FSLS method uses an interval time t for flooding, and sets the TTL value to S_i such that S_i appears at every $2^{(i-1)} \cdot t$ interval. The FSLS method delivers information at short intervals to the neighboring areas of a node that has had a change in link utilization, and delivers information at longer intervals to more distant (in terms of hop–count) nodes. Note that if multiple changes in link utilization occur, the FSLS method aggregates information, which decreases the number of control messages. Ref. [7] evaluated the performance of the FSLS method mainly in a random network, but the number of control messages in a power–law network was not described.

Other reports [6, 8, 9] propose flooding methods based on information retrieval in P2P (Peer–to–Peer) networks, and evaluate the amount of search queries in a power–law network. Information retrieval from a peer is performed by forwarding a search query to other peers. One issue highlighted in these studies is that for efficient information retrieval, the amount of query forwarding has to be reduced, while each search query has to reach as many peers as possible. This is similar to our problem in that route information must be delivered with less forwarding, i.e., with fewer control messages. The flooding method for a P2P network does not require the delivery of a query to all nodes. Based on this observation, a probabilistic relaying method has been proposed [6]. Percolation theory [5] is used to obtain the relay probability. However, with our method of distributing routing information, the information has to be finally acquired by all nodes, otherwise route mismatching could occur at individual nodes, causing problems with route convergence [10].

III. Evaluation of Simple Flooding Method in Power–Law Networks

In this section, we evaluate the number of control messages generated by flooding in power–law and random networks. Then, we examine the problems with a simple flooding method that relays messages to all neighboring nodes in a power–law network.

A. Simulation Model

For the network model, we used a 1000–node network topology generated by the BA model [4], which is one of the power–law network generating models. In the BA model, the number of initial nodes m_0 and the number of additional links per node m are set to 2, respectively. We also used a 1000–node network topology generated by the ER model [11] as a random network model. In the ER model, the probability

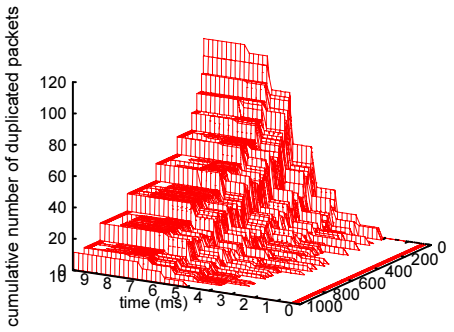

(a) Cumulative number of duplicated messages generated by single–node failure in 1000–node random network

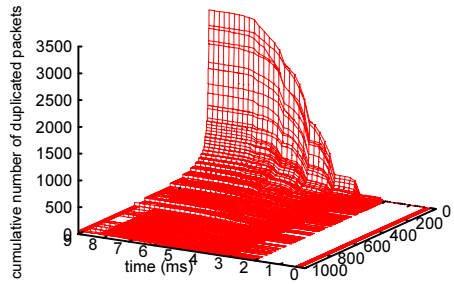

(b) Cumulative number of duplicated messages generated by hub–node failure in 1000–node power–law network

Fig. 1. Number of duplicated messages generated by single–node failure in 1000–node network

of connecting with each node is set at $1/(N-1)$. Here, N is the number of network nodes. This probability generates almost the same number of links as the BA model under the same number of nodes. Only the degree distribution is different between the topology produced by the BA model and that by the ER model. Other parameters were as follows. The propagation delay on each link was set at 1 ms. There was no communication traffic except control messages. Each node process packet was based on an FIFO queue and the packet processing capability was set at 1Mpps.

We performed flooding from multiple nodes assuming node failure, and we evaluated the number of control messages generated in each network.

B. Evaluations

Figure 1 shows the cumulative number of control messages that each node received depending on the time after a node broke down. The y–axis shows the node indexes in ascending order of the number of links they connected with. Figure 1(b) shows that, in power–law networks, there is little duplication of control messages at nodes with a small number of links. However, it also shows that a hub–node that has a large number of links receives an enormous number of control messages. As Fig. 1(a) shows, there are few duplicated control messages in random networks compared to in power–law networks.

Next, in Fig. 3, we show the maximum number of duplicated messages that arrived at a node when the number of nodes in both networks was changed. Note that the results are averaged

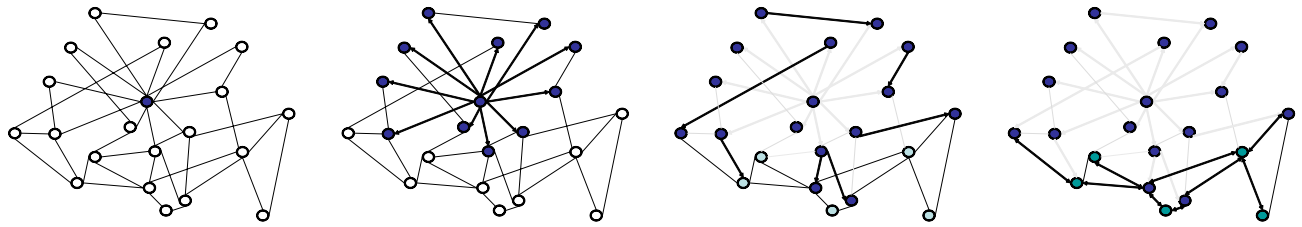

(a) Flooding occurs from center node (b) Message delivered to neighbor nodes with high relay probability (c) Each node sends message to its own neighbor nodes with low relay probability (d) Any nodes that are missed originally receive message via periodical exchanges

Fig. 2. Mechanism of proposed flooding method

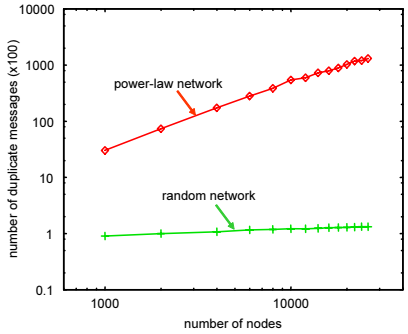

Fig. 3. Relationship between number of duplicated messages with increase in number of nodes

over ten experiments. As this figure shows, the number of duplicated messages increased rapidly as the number of nodes increased in the power–law network, whereas it increased slowly in the random network. Even in networks consisting of 26,000 nodes, there were only around 130 duplicated messages in the random network, while there were 130,000 duplicate messages in the power–law network. This clearly indicates that congestion is more likely to occur in a power–law network than in a random network, and that a power–law network does not scale well due to the concentration of messages at hub–nodes.

IV. Proposed Method

The FSLS method reduces the number of control messages in a network. Nodes that are further away (in terms of hop–counts) from the node that originates a flood have less opportunity to receive control messages. Furthermore, the nodes receive aggregated information of link status changes that occur at a $2^{(i-1)} \cdot t$ interval. However, since the FSLS sets a low TTL initially, flooding from non–hub nodes takes a long time to deliver route information. Flooding using probabilistic relaying, which we call probabilistic flooding, reduces the number of duplicated control messages. However, by its very nature, this method cannot guarantee to deliver route information to the entire network, i.e., there is a possibility that a node will not receive information from neighbor nodes.

To solve this problem in the probabilistic flooding method, this paper proposes a new flooding method for power–law networks. Our method has the three following features:

i) The number of control messages in a network is reduced by probabilistic relaying,
ii) To guarantee that information is delivered to all nodes, each node exchanges information periodically,
iii) The relaying probability at a node changes according to whether it is the node originating a flood.

The last feature (iii) increases the number of nodes that receive route information. Determination of the relaying probability, which is an important issue for our method, is described in the next section.

Figure 2 shows an example of the proposed method. When the center node initiates a flood (Fig. 2(a)), the node relays a message to the next node with a high relaying probability (Fig. 2(b)). Nodes that receive the message relay it to their neighbor nodes with a lower relaying probability, as shown in Fig. 2(c). As mentioned above, probabilistic flooding sometimes fails to deliver information to all nodes. Then, each node exchanges information with its adjacent nodes to ensure information is delivered to any nodes that did not receive the information originally by probabilistic flooding (Fig. 2(d)).

Information exchange is processed by the way that, (i) a node probes the information of neighboring nodes periodically, (ii) and receives the new information if the neighbor have new information. The overhead of information exchanging can be ignored because we can include the probing process into the neighbor–finding messages in a routing protocol (like HELLO messages in OSPF protocol).

A. Choosing the Relay Probability

In a previous study [5] a relaying probability was derived in which almost all nodes received messages. Percolation theory was applied to obtain the probability mathematically. Since the γ of the topology we discussed in Section III was 2.39 and the maximum degree in the network was 53, the relay probability p_c was set at 0.9 based on this study [5].

Figure 4 shows the information reachability, which is defined as the ratio of nodes that receive the information to the nodes in the network, and the number of control messages. The figure shows the results for different relay probabilities p. Looking at Fig. 4, setting the relay probability at 0.9 ensures that messages are delivered to almost all nodes. In this case, the number of control messages in the network was 10% less than with simple flooding (See Fig. 4(b)).

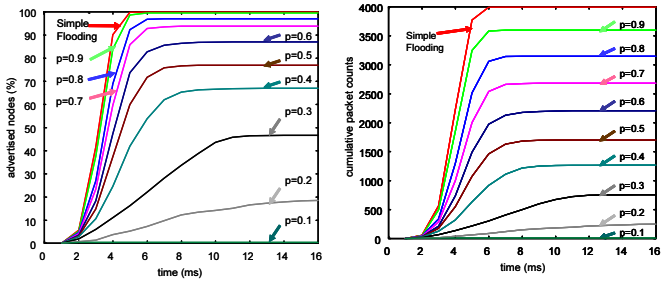

(a) Relationship between relay probability and information reachability (b) Relationship between relay probability and number of control messages

Fig. 4. Simple probabilistic method from nodes with a large number of links

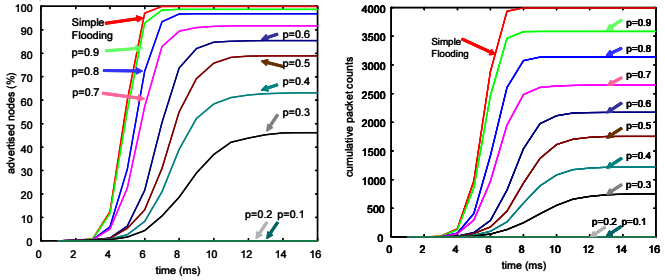

(a) Relationship between relay probability and information reachability (b) Relationship between relay probability and number of control messages

Fig. 5. Simple probabilistic method from nodes with a small number of links

In our method, each node performs periodical information exchange to guarantee information delivery to all nodes. Hence, it is not necessary to set the relaying probability at p_c. Rather, we use a lower relay probability in the expectation that there will be fewer control messages. In this paper, we set the relay probability using an empirical approach. The simulation results, which are not presented here, indicated that periodical information exchange delivered route information to around 15% of nodes. We therefore chose a relay probability at which 85% of nodes received route information. That is, the relay probability was set at 0.6 according to Fig. 4.

B. Flooding from Non–Hub Nodes

With a lower relay probability, flooding from non–hub nodes becomes a problem. Figure 5(a) shows a typical case of flooding from non–hub nodes. As shown in the figure, information may not be spread throughout the network even when the relay probability is set at $p_c(= 0.9)$. We therefore introduced two types of relay probability: p_1 and p_2. The p_1 is the relay probability from a node that starts a flood and the p_2 is the probability with which neighbor nodes relay the message to their corresponding neighbor nodes. In this paper, we set p_1 at 1.0 so that information would be widely delivered by flooding from non–hub nodes, ensuring that adjacent nodes would receive the information. p_2 was set at 0.6, as described in the previous section.

V. Evaluation of the Proposed Method

In this section, we describe the results of a simulation to evaluate the validity and effectiveness of the proposed method. As a network model, we used the same topology used in section III, i.e., a power–law topology with 1,000 nodes. This method was compared with three conventional methods: a simple flooding method, which delivers information to all nodes; a probabilistic flooding method, in which the relay probability is chosen on the basis of percolation theory [6]; and the FSLS method. In the probabilistic flooding method, the relay probability p_c was set at 0.9 as described in section IV, since the γ of the topology used for the simulation was 2.39 and the cutoff K was 53. In the FSLS method, we set $S_i = i$ $(1 \leq i \leq 7)$ and the interval for changing the TTL at 1 second.

Figure 6(a) shows the number of nodes that received information when a hub–node performed flooding and Figure 6(b) shows the cumulative number of control messages. As these figures show, neither the simple flooding method nor the probabilistic flooding method decreased the number of control messages. In contrast, our proposed method reduced the number of control messages by about 60%, while around 85% of the nodes in the network received control messages. The remaining 15% of nodes received route information via an exchange of information between neighboring nodes after a fixed period (here, we assumed a 5–second period). When flooding from non–hub nodes occurred, the number of messages was reduced by about 50%, but these figures are omitted due to space limitations.

The results for the FSLS method indicated that route information spread slowly since the FSLS set a low TTL initially. Furthermore, the FSLS method generated more control messages with time. This is because the nodes nearest the node that starts a flood receive the route information more than once. For example, nodes that receive route information when the TTL is S_2 also receive the information when the TTL is S_3, S_4, S_5, or S_6.

Figure 7 shows the information reachability and number of control messages generated when multiple nodes perform flooding. Here, node failure occurs based on a Poisson arrival process with a mean arrival rate of 1 / second. The node is selected randomly. The proposed method reduced the number of control messages in the network by 50% compared with the simple flooding method. We also found that the proposed method decreased the number of control messages by about 40% compared with the probabilistic flooding method.

We also examined these four methods on a power-law topology with 3,000 nodes generated by BA model, and got almost the same results as above. This shows that our method performs the same in larger networks.

VI. Conclusion

In this paper, we evaluated the performance of flooding mechanisms in a power–law network. The results showed that congestion was more likely to occur in a power–law network than in a random network, and that power–law networks did not scale well due to the concentration of messages at hub–nodes. We therefore proposed a new flooding method for power–law networks. The proposed method reduces the

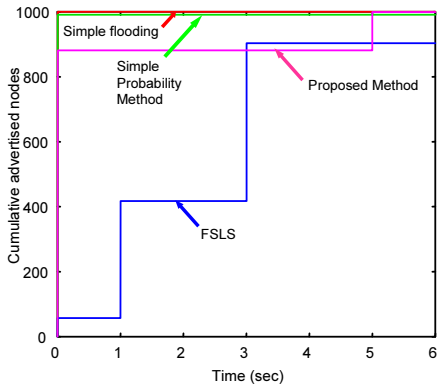

(a) Number of nodes that receive information

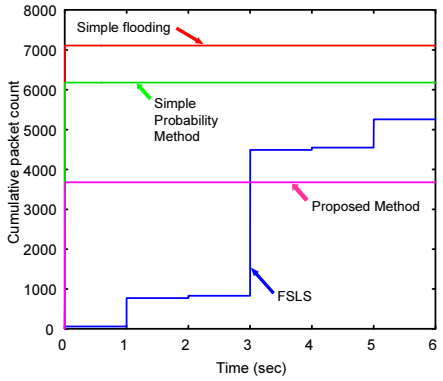

(b) Cumulative number of control messages

Fig. 6. Evaluation of number of control messages with proposed method

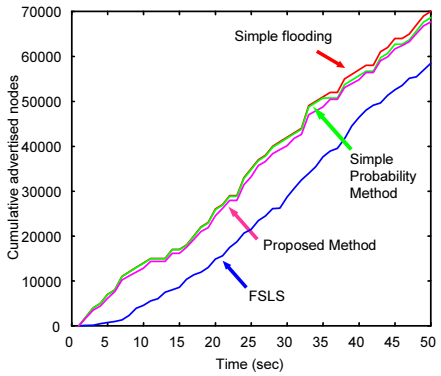

(a) Cumulative number of nodes that received information

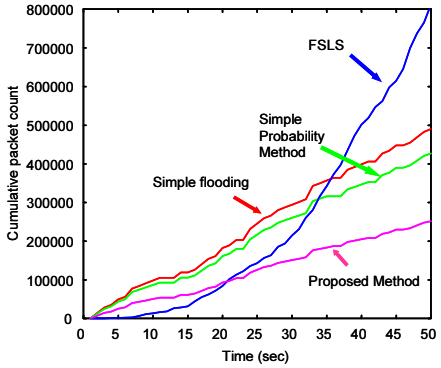

(b) Cumulative number of control messages

Fig. 7. Evaluation of proposed method with multiple flooding

duplication of control messages by using probabilistic relaying of messages, but guarantees that information is delivered to all nodes by exchanging information between adjacent nodes periodically. Our proposed method can also use a lower relaying probability than that of conventional methods because of this periodic exchange of information between nodes.

Simulation results showed that the proposed method decreased the number of control messages by about 50% compared with a simple flooding method, and by about 40% compared with a conventional probabilistic flooding method.

In future, we plan to first derive the relaying probability of the proposed method theoretically, although it was chosen according to the results of the simulation. Secondly, we used the BA model to generate the model of a power–law network. However, other studies (e.g., Ref. [12]) propose different models for constructing power–law networks, which indicate that the topology produced by the BA model does not reflects the topology characteristics of the Internet. We will therefore evaluate flooding mechanisms using different network models.

REFERENCES

[1] A. Shaikh, C. Isett, A. Greenberg, M. Roughan, and J. Gottlieb, "A case study of OSPF behavior in a large enterprise network," in *Proceedings of ACM SIGCOMM Internet Measurement Workshop (IMW)*, pp. 217–230, Nov. 2002.

[2] L. Gao and J. Rexford, "Stable Internet routing without global coordination," in *Proceedings of ACM SIGMETRICS 2000*, pp. 307–317, June 2000.

[3] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power–law relationships of the Internet topology," in *Proceedings of ACM SIGCOMM '99*, pp. 251–262, Oct. 1999.

[4] A. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, pp. 509–512, Oct. 1999.

[5] R. Cohen, S. Havlin, and D. Avraham, "Structural properties of scale–free networks," in *Handbook of Graphs and Networks – From the Genome to the Internet* (S. Bornholdt and H. G. Schuster, eds.), WILEY-VCH GmbH & Co., 2003.

[6] F. Banaei-Kashani and C. Shahabi, "Criticality–based analysis and design of unstructured peer-to-peer networks as 'complex systems'," in *Proceedings of Third International Workshop on Global and Peer-to-Peer Computing (GP2PC)*, pp. 22–32, May 2003.

[7] C. Santivañez, R. Ramanathan, and I. StavrakakisPablo, "Making link-state routing scale for ad hoc networks," in *Proceedings of ACM Mobihoc 2001*, pp. 22–32, Oct. 2001.

[8] R. Guimerà, A. Díaz-Guilera, F. Vega-Redondo, A. Cabrales, and A. Arenas, "Optimal network topologies for local search with congestion," *Physical Review Letters*, vol. 89, 248701, Dec. 2002.

[9] G. M. Viswanathan, S. V. Buldyrev, S. Havlln, M. G. E. da Luz, E. P. Raposo, and H. E. Stanley, "Optimizing the success of random searches," *Nature*, vol. 401, pp. 911–914, Oct. 1999.

[10] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed internet routing convergence," in *Proceedings of SIGCOMM*, pp. 175–187, Aug. 2000.

[11] M. E. J. Newman, "Random graphs as models of networks," in *Handbook of Graphs and Networks – From the Genome to the Internet* (S. Bornholdt and H. G. Schuster, eds.), Berlin: WILEY-VCH GmbH & Co., 2003.

[12] L. Li, D. Alderson, W. Willinger, and J. Doyle, "A first-principles approach to understanding the Internet's router-level topology," in *Proceedings of ACM SIGCOMM 2004*, pp. 3–14, ACM Press, 2004.