On the Effect of Scale-Free Structure of Network Topology on TCP Performance

Hiroyuki Ohsaki and Makoto Imase
Graduate School of Information Science and Technology, Osaka University, Japan
1-5 Yamadaoka, Suita, Osaka 565-0871, Japan
Tel: +81-6-6879-4551, Fax: +81-6-6879-4554
Email: {oosaki,imase}@ist.osaka-u.ac.jp

***Abstract*— In recent years, it has been reported that several existing networks represented by the Internet have scale-free structure. In this paper, through simulation experiments, we investigate the effect of the scale-free structure of communication networks on the end-to-end performance of TCP flows. As network topology, a random network and a scale-free network with the equal number of nodes and the equal number of links are used. We compare the end-to-end performance of TCP flows (i.e., throughput, round-trip time, and packet loss rate) in a random network and a scale-free network. Consequently, we show that, contrary to common beliefs, the scale-free structure of a network is someties halmful on the network performance. Namely, we show that, although the scale-free structure of a network has positive effect on the round-trip time of TCP flows, it has negative effect on the throughput and the packet loss rate of TCP flows when the average degree of a network is large.**

I. INTRODUCTION

In recent years, attention to a topology of large-scale communication networks has been increasing [1]. Studies on a topology of communication networks have long history [2, 3], but those conventional researches have focused on comparatively small-scale communication networks. In the late 1990s, it was discovered that several real networks such as the topology of Internet ASs and the hyperlink structure of Web pages exhibit scale-free structure [4].

Such finding causes increasing concern on the optimal topology of, in particular, large-scale communication networks. Most conventional researches on a topology of communication networks have generally focused on regular networks such as star, ring, and mesh, and random networks. However, as several interesting characteristics of a scale-free network become clear, the relation between a communication network and its scale-free structure has been attracting much attention.

Notable characteristics of a scale-free network include, for instance, that the average distance of a network (i.e., the average of shortest path lengths between arbitrary node pair) is much smaller than that of a random network, and that a scale-free network is more robust to random node failures (i.e., connectivity among nodes is more likely to be preserved). Hence, researches on a topology of large-scale communication networks for improving reliability [5] and/or for improving packet transfer efficiency [6, 7] have been performed.

In recent years, there exist researches focusing not only on network-level performance (e.g., reliability, resilience, and per-link transfer efficiency), but also on user-level performance (i.e., end-to-end performance), which should be the most important metrics to users.

The characteristic of a scale-free network that the average distance of a network is small is advantageous for performing information retrieval and maintaining reachability. However, considering packet transfer over a communication network, such a small average distance of a scale-free network implies traffic concentration at hub nodes. If traffic is concentrated at hub nodes, those nodes would become the bottleneck of the network, and limit the performance of the entire network. Namely, from a viewpoint of the end-to-end performance, a small average distance of a scale-free network and traffic concentration at hub nodes should have opposite effects on the end-to-end performance.

In [8], through a simple numerical analysis, we have clarified the effect of the scale-free structure of a network on its end-to-end performance. In [8], the average throughput of each flow was derived by assuming that the bandwidth allocation to each flow satisfies Max-Min fairness. Consequently, we have found: (1) when the average degree of a network is small (i.e., there are not many links in a network), a scale-free network shows higher end-to-end performance than a random network, and (2) conversely, when the average degree of a network is large (i.e., many links in a network), a random network shows higher end-to-end performance than a scale-free network.

In these days, majority of the Internet traffic has been carried by TCP (Transmission Control Protocol) [9]. It is known that the bandwidth allocation to TCP flows satisfies proportional fairness, instead of Max-Min fairness, since TCP has the window-based flow control and packet retransmission mechanism [10]. Also, it is known that the end-to-end performance of TCP flows will deteriorate greatly when a large number of TCP flows compete the shared network resource. Hence, for investigating the effect of the scale-free structure of a network on the end-to-end performance, it is necessary to take account of the effect of such end-to-end congestion control mechanisms.

Moreover, not only throughput of TCP flows but also round-trip time and packet loss rate of TCP flows are important performance metrics to users. So, it is important to investigate the effect of the scale-free structure of a network on those end-to-end performance metrics.

In this paper, we therefore investigate the effect of the scale-free structure of a network on the end-to-end performance of TCP flows through simulation experiments. Similar to [8], as network topology, a random network and a scale-free network with the equal number of nodes and the equal number of links are used. We compare the end-to-end performance of TCP flows (i.e., throughput, round-trip time, and packet loss rate) in a random network and a scale-free network. Consequently, we show that, contrary to common beliefs, the scale-free structure of a network is someties halmful on the network performance. Namely, we show that (1) the scale-free structure of a network has positive effect on the round-trip time of TCP flows regardless of the number of nodes and the average degree of a network, and (2) the scale-free structure of a network has negative effect on the throughput and the packet loss rate of TCP flows when the average degree of a network is large.

The organization of this paper is as follows. First, related works are explained in Section II. Section III briefly summarizes fundamental characteristics of a random network and a scale-free network. In Section IV, effect of the scale-free structure of a network on the end-to-end performance of TCP flows is evaluated quantitatively through simulation experiment. Finally, in Section V, we conclude this paper and discuss future works.

II. RELATED WORKS

In [11], the effect of a network topology on router load (i.e., the number of packets processed at a router per unit time) is investigated. The authors of [11] assume a simple network model where (1) each node transmits a packet to all other nodes at a fixed rate, and (2) routing is determined simply by the number of hops (i.e., packets traverse the shortest path among possible routes). The authors reveal existence of correlation between the number of packets processed at a router and the degree of the router. Also, the authors show that such correlation becomes strong as the amount of traffic in a network increases. However, the authors of [11] focus only on router load, which is one of network-level performance metrics, and does not take account of the end-to-end performance. Also, they do not take account of the effect of TCP congestion control.

Also, in [12], the effect of a network topology on router load is investigated. The authors of [12] investigate the distribution of router load in a random network and a scale-free network. The authors assume a simple network model where (1) a node generates a packet following a distribution defined by an LRD (Long-Range Dependence) model, (2) a packet moves to the neighboring node at every unit time, and (3) routing is determined simply by the number of hops (i.e., packets traverse the shortest path among possible paths). The authors of [12] show that the communication performance of a scale-free network is lower than that of a random network. However, since the authors assume a quite simplified network model, it is difficult to generalize this result to other packet switching networks.

III. RANDOM NETWORK AND SCALE-FREE NETWORK

In this section, fundamental characteristics of a random network and a scale-free network are summarized. Moreover, models that generate either a random network or a scale-free network are also explained.

A random network is a network where the probability that a link exists between arbitrary node pair is given by a uniform distribution [13].

As a representative model for generating a random network, ER (Erdos-Renyi) model [14] has been widely used. Two parameters, the number N of nodes and the connection probability p between nodes, are used for generating a random network. For a given N nodes, a random network is generated by creating links among all node pairs with the probability p.

As characteristics of a random network, it is known that the degree distribution $P(k)$ follows a binomial distribution and, for sufficiently large N, the average distance l satisfies $l \propto \log N$ [15].

A scale-free network is a network where its degree distribution follows the following power-law [4].

$$P(k) \propto k^{-\lambda} \tag{1}$$

In the literature, several models that generate scale-free networks have been proposed [16, 17]. In this paper, we explain the BA (Barabasi Albert) model [18], which is one of the most representative models for generating a scale-free network. Notable features of the BA model are network growth and link preferential attachment [18]. First, a connected network with a small number of nodes is created, and nodes are added to the network one-by-one. Then, a scale-free network can be generated by creating a link between a new node and existing one, which is randomly chosen from all existing nodes with a probability proportional to the degree of an existing node. It is known that the network generated by the BA model has a power-law index of $\lambda = 3$.

As characteristics of a scale-free network, it is known that the average distance l is much smaller than that of a random network. For instance, it is known that, for a sufficiently large N, the average distance satisfies $l \propto \log\log N$ for $2 < \lambda < 3$ [15].

IV. SIMULATION

A. Model

In what follows, we investigate the effect of the scale-free structure of a network on the end-to-end performance of TCP flows by simulation experiments.

As network topologies, a random network and a scale-free network (generated by the BA model) are used. The number of nodes (i.e., routers or hosts) in a network is denoted by N, and the average degree (i.e., the average number of links connected to a node) is denoted by k. We compare the end-to-end performance of TCP flows in a random network and a scale-free network with the same number of nodes N and the average degree k. Thereby, we clarify the effect of the scale-free structure of a network on the end-to-end performance of TCP flows.

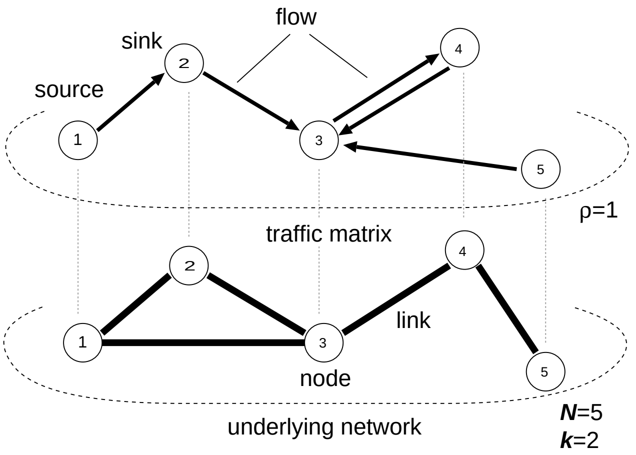

Fig. 1: Flow model and underlying network model

On the random network and the scale-free network, TCP flows were generated randomly (see Fig. 1). Specifically, TCP flows on each network was given by another random network with N nodes and the average degree $\rho \times k$ (hereafter, ρ is called *load factor*). Namely, TCP flows were generated from node i to node j, when there exist a link from node i to node j in the random network. Since the number of links in a network is $k \times N/2$, load factor ρ is a parameter that determines the ratio of the number of flows to the number of links in a network.

We generated 20 random networks and 20 scale-free networks. As metrics of the end-to-end performance, we measured throughput, round-trip time, and packet loss rate of each TCP flow in both networks. Moreover, we measured the number of TCP flows traversing a router.

Unless explicitly stated, the parameters shown in Tab. I are used in the following simulations. For simplicity, all link have the same bandwidth B, and all links have the same propagation delay τ. In the following simulation results, 95% confidence interval is plotted. ns-2 (version 2.29) were used for simulation. Note that routing of TCP flows is determined by the shortest path algorithm.

TABLE I

PARAMETER CONFIGURATION USED IN SIMULATION

number of nodes	N	100	
average degree	k	3	
link bandwidth	B	10	[Mbit/s]
propagation delay	τ	10	[ms]
router buffer size	L	100	[packet]
load factor	ρ	5.0	

B. Number of TCP Flows Traversing a Router

First, we focus on the number of TCP flows traversing a router in a random network and a scale-free network. Thereby, we investigate how equally every router is utilized by TCP flows.

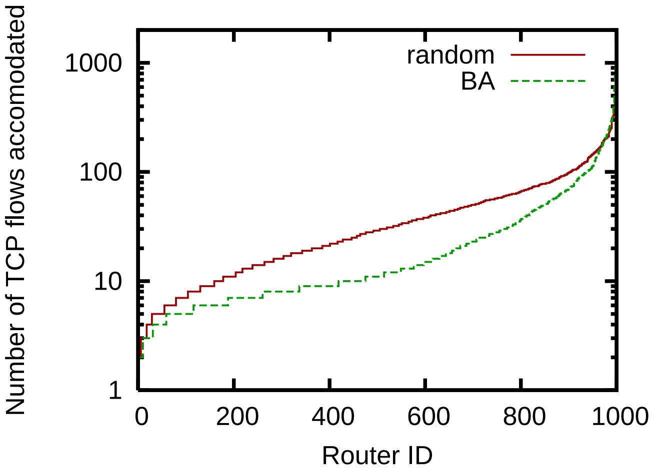

Fig. 2: Distribution of the number of TCP flows traversing a router ($N = 1,000$, $k = 3$)

The distribution of the number of TCP flows traversing a router in a random network and a scale-free network for $N = 1,000$ nodes and the average degree of $k = 3$ is shown in Fig. 3.

This figure shows that the number of TCP flows traversing a router in a random network is larger than in a scale-free network. We calculated the average number of TCP flows traversing a router: 47.75 in the random network and 33.96 in the scale-free network. However, this figure also shows that in a scale-free network, there exist a small number of routers accommodating a large number of TCP flows. Namely, this means that *hub routers* accommodate a large number of TCP flows.

Thus, in a random network, the deviation of the number of TCP flows traversing a router is relatively small, and all routers are comparatively equally utilized. On the contrary, in a scale-free network, the deviation of the number of TCP flows traversing a router is relatively large, and some hub routers are heavily utilized. This observation corresponds to the findings in [19, 12].

In what follows, we investigate the effect of such characteristics of a scale-free network on the end-to-end performance of TCP flows.

C. Average TCP Flow Throughput

First, we focus on the throughput of TCP flows as one of the end-to-end performance metrics.

The average throughput of TCP flows when fixing the average degree at $k = 3$ and changing the number N of nodes from 100 to 1,000 is shown in Fig. 3. This figure shows that the throughput of the scale-free network is approximately 10% larger than that of the random network regardless of the number of nodes.

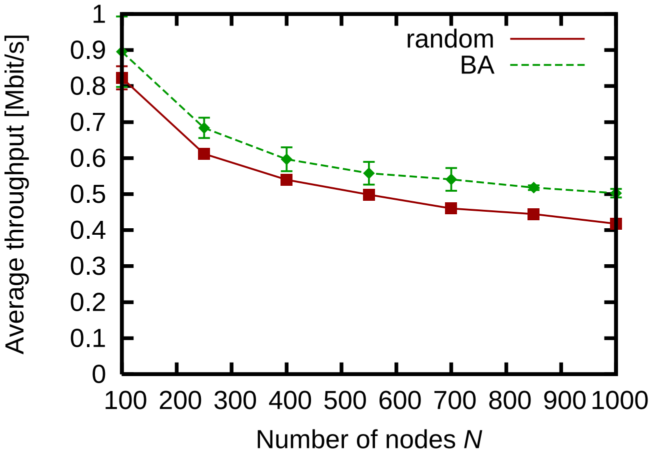

Fig. 3: Average TCP flow throughput for different number of nodes N ($k = 3$, $B = 10$ [Mbit/s])

As discussed in Section I, from a viewpoint of the end-to-end performance, a small average distance of a scale-free network and traffic concentration at hub routers should have opposite effects on the end-to-end performance. In Fig. 3, the positive effect, such that the average distance of a scale-free network is small, should have stronger effect than the negative effect.

However, when the average degree of a network is changed, we can observe quite different tendency. The average throughput of TCP flows when fixing the number of nodes at $N = 1,000$ and changing the average degree k from 2 to 8 is shown in Fig. 4. This figure shows that the random network shows larger throughput than the scale-free network when the average degree k is large. Specifically, the relation of average throughputs in a random network and a scale-free network is reversed at approximately $k = 5$.

This is probably because the negative effect, such that traffic is likely to be concentrated on hub routers, is significant when the average degree is large. Although results are not included in this paper due to space limitation, we found that the relation of the average throughput in a random network and a scale-free network is reversed at approximately $k = 6$ when the average throughput is calculated as in [8] (not modeling TCP congestion control mechanism). This result indicates that the negative effect, such that traffic is likely to be concentrated on hub routers, is significant when the TCP performs its congestion control.

From these observations, we conclude that the scale-free structure of a network has positive effect on the throughput of TCP flows when the average degree of a network is small, but it has negative effect on the throughput of TCP flows when the average degree of a network is large.

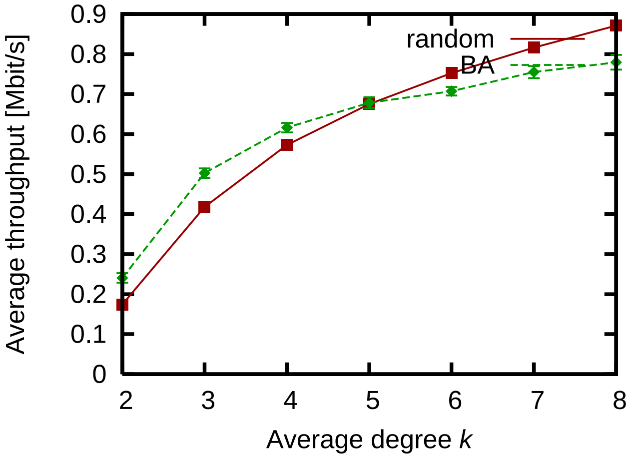

Fig. 4: Average TCP flow throughput for the different average degree k ($N = 1,000$, $B = 10$ [Mbit/s])

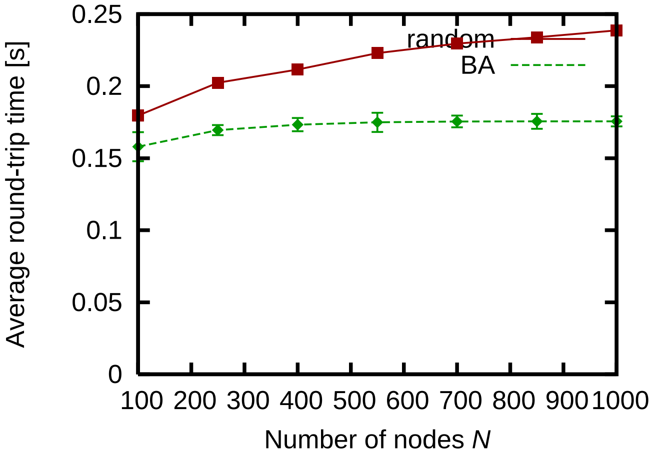

Fig. 5: Average round-trip time for different number of nodes N ($k = 3$, $B = 10$ [Mbit/s])

D. Average Round-Trip Time

Next, we focus on the round-trip time of TCP flows as one of the end-to-end performance metrics.

The average round-trip time of TCP flows (i.e., the average of round-trip times of all TCP flows) when fixing the average degree at $k = 3$ and changing the number N of nodes from 100 to 1,000 is shown in Fig. 5. This figure shows that the scale-free network has smaller average round-trip time than the random network regardless of the number N of nodes. Also, this figure shows the difference in average round-trip times becomes large as the number N of nodes becomes large. This is probably caused by a small average distance of a scale-free network as discussed in Section III.

It should be noted that this tendency does not change so

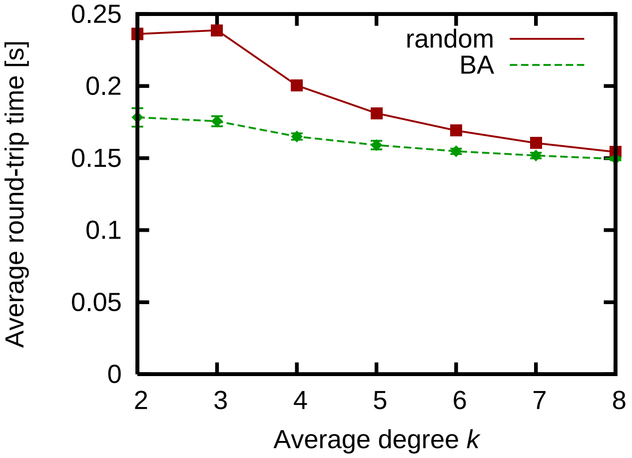

Fig. 6: Average round-trip time for different average degree k ($N = 1,000$, $B = 10$ [Mbit/s])

much even when the average degree k is changed. The average round-trip time of TCP flows when fixing the number of nodes at $N = 1,000$ and changing the average degree k from 2 to 8 is shown in Fig. 6. This figure shows that the scale-free network has smaller average round-trip time than the random network regardless of the average degree k. Although the difference becomes small as the average degree k becomes large, the relation is not reversed upto the average degree of $k = 8$.

By comparing the average TCP throughput (Fig. 4) and the average round-trip time (Fig. 6), we can see interesting phenomenon.

When the average degree k is large, the average round-trip time of the random network is larger than that of the scale-free network (Fig. 6). Generally, it is known that the throughput of TCP flows is inverse proportional to its round-trip time [20]. However, the average throughput of the random network is larger than that of the scale-free network (Fig. 4). This phenomenon can be explained as follows. In a scale-free network, TCP flows are likely to be concentrated on hub routers, so that the throughput of TCP flows in the scale-free network degrades severely.

From these observations, we conclude that the scale-free structure of a network has positive effect on the round-trip time of TCP flows regardless of the number of nodes and the average degree.

E. Packet Loss Rate

Finally, we focus on the packet loss rate of TCP flows as one of the end-to-end performance metrics.

The packet loss rate of TCP flows (i.e., the average of packet loss rates of all TCP flows) when fixing the average degree at $k = 3$ and changing the number N of nodes from 100 to 1,000 is shown in Fig. 7. This figure shows that the packet loss rate of TCP flows in a random network and a scale-free network are almost same regardless of the number of nodes.

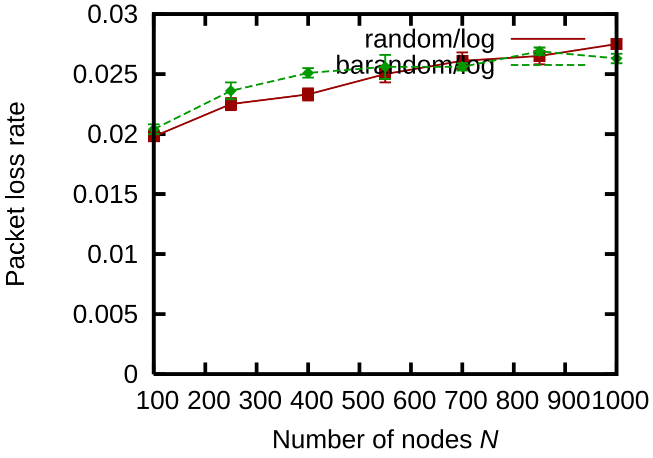

Fig. 7: Packet loss rate for different number of nodes N ($k = 3$, $B = 10$ [Mbit/s])

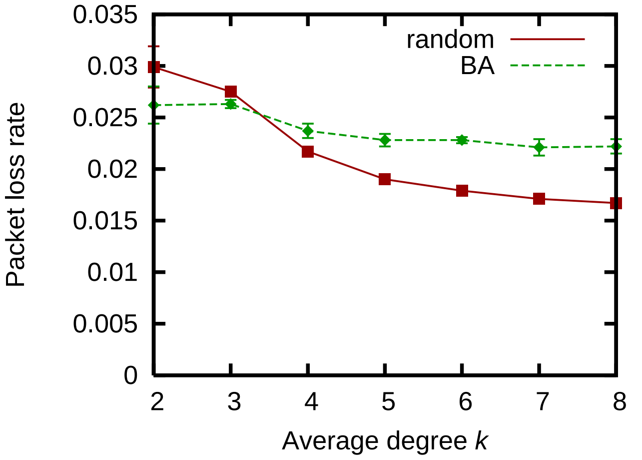

Fig. 8: Packet loss rate for different average degree k ($N = 1,000$, $B = 10$ [Mbit/s])

The packet loss rate of TCP flows when fixing the number of nodes at $N = 1,000$ and changing the average degree k from 2 to 8 is shown in Fig. 8. This figure shows that the relation between the packet loss rates in a random network and a scale-free network is determined by the average degree k.

In this figure, the packet loss rate of TCP flows in a scale-free network does not change so much even when the average degree k is changed. This is because TCP flows are likely to be concentrated on hub routers in a scale-free network regardless of its average degree.

On the other hand, the packet loss rate of TCP flows in a random network decreases as the average degree k becomes large. This is because TCP flows in a random network are likely to be distributed to multiple routers when the average degree k is large.

From these observations, we conclude that the scale-free structure of a network has positive effect on the packet loss rate of TCP flows when the average degree of a network is small, but it has negative effect on the packet loss rate of TCP flows when the average degree of a network is large.

F. Discussion

Our simulation results clearly indicate that the scale-free structure of a network has positive effect on the round-trip time of TCP flows regardless of the number of nodes and the average degree. Namely, even when the effect of the TCP congestion control is taken account of, the scale-free structure of a network has positive effect for delay-sensitive applications (e.g., information retrieval and realtime communication). Moreover, it should be noted that such a characteristic of a scale-free network is independent of the number of nodes and the average degree of the network.

On the other hand, the scale-free structure of a network has negative effect on the throughput and packet loss rate of TCP flows when the average degree of a network is large. Namely, because of characteristics of a packet switching network and the TCP congestion control mechanism, the scale-free structure of a network has negative effect for non-delay-sensitive applications (e.g., file transfer and Web browsing).

When the average degree of a network is small, the scale-free structure of a network has positive effect on the throughput and packet loss rate of TCP flows. However, considering that the average degree of typical networks ranges from 4.7 to 6.3 [21], we believe that the scale-free structure of a network generally has negative effect for bandwidth-intensive applications.

V. CONCLUSION

In this paper, through simulation experiments, we have investigated effect of the scale-free structure of communication networks on the end-to-end performance of TCP flows. Consequently, we have shown that, contrary to common beliefs, the scale-free structure of a network is someties halmful on the network performance. Namely, we have shown that (1) the scale-free structure of a network has positive effect on the round-trip time of TCP flows regardless of the number of nodes and the average degree of a network, and (2) the scale-free structure of a network has negative effect on the throughput and the packet loss rate of TCP flows when the average degree of a network is large.

As future work, we are planning to mathematically analyze the relation between the scale-free structure of a network and the end-to-end performance of TCP flows. Also, we are planning to investigate the effect of the scale-free structure of a network on non-standard TCP protocols and UDP-based transport protocols.

ACKNOWLEDGMENTS

We would like to express our appreciation to Mr. Koutaro Yagi of the Graduate School of Information Science and Technology, Osaka University, Japan for joining meaningful discussions and theoretical arguments.

REFERENCES

[1] R. Albert and A. L. Barabasi, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, no. 1, pp. 48–98, Jan. 2002.

[2] L. Hu, "Topology control for multihop packet radio networks," *IEEE Transactions on Communications*, vol. 41, pp. 1474–1481, Oct. 1993.

[3] R. Albert, H. Jeong, and A. L. Barabsi, "Diameter of the world wide web," *Nature*, vol. 401, pp. 130–131, Sept. 1999.

[4] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," *ACM SIGCOMM Computer Communication Review*, vol. 29, pp. 251–263, May 1999.

[5] Y. Tu, "How robust is the Internet?" *Nature*, vol. 406, pp. 353–354, May 2000.

[6] B. Tadic and G. J. Rodgers, "Packet transport on scale free networks," *Advances in Complex Systems*, vol. 5, pp. 445–456, July 2002.

[7] C. M. Ghim, E. Oh, K. I. Goh, B. Kahng, and D. Kim, "Packet transport along the shortest pathways in scale-free networks," *The European Physical Journal B*, vol. 38, pp. 193–199, Feb. 2004.

[8] H. Ohsaki, K. Yagi, and M. Imase, "On the effect of scale-free structure of network topology on end-to-end performance," in *Proceedings of the 2007 International Symposium on Applications and the Internet (SAINT 2007)*, Jan. 2007, pp. 12–18.

[9] V. Paxson, "Growth trends in wide-area TCP connections," *IEEE Network*, vol. 8, no. 4, pp. 8–17, July 1994.

[10] M. Vojnovid, J.-Y. L. Boudec, and C. Boutremans, "Global fairness of additive-increase and multiplicative-decrease with heterogeneous round-trip times," in *Proceedings of IEEE INFOCOM*, 2000, pp. 26–30.

[11] G. Peli and G. Papp, "Are scale free networks better?" http://arxiv.org/pdf/cond-mat/0301555, Feb. 2003.

[12] D. Arrowsmith, M. Bernardo, and F. Sorrentino, "Effects of variations of load distribution on network performance," in *Proceedings of IEEE Systems and Circuits*, May 2005, pp. 3773–3776.

[13] B. Bloolbas, *Random Graphs Second Edition*. Cambridge Univ Press, Oct. 2001.

[14] P. Erdos and A. Renyi, "On random graphs," *Mathematica*, vol. 6, pp. 290–297, Dec. 1959.

[15] R. Cohen and S. Havlin, "Scale-free network are ultra small," *Physical Review Letters*, vol. 90, p. 058701, Feb. 2003.

[16] G. Bianconi and A. L. Barabasi, "Competition and multiscaling in evolving networks," *Europhysics Letters*, vol. 54, pp. 436–442, May 2001.

[17] K. Klemm and V. M. Eguiluz, "Highly clustered scale-free networks," *Physical Review E*, vol. 65, p. 036123, Feb. 2002.

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

[19] K. I. Goh, B. Kahng, and D. Kim, "Packet transport and load distribution in scale-free network models," *Physica A*, vol. 318, pp. 72–79, Feb. 2003.

[20] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP throughput: a simple model and its empirical validation," in *Proceedings of ACM SIGCOMM '98*, Sept. 1998, pp. 303–314.

[21] P. Mahadevan *et al.*, "The internet AS-level topology: three data sources and one definitive metric," *ACM SIGCOMM Computer Communication Review*, vol. 36, pp. 17–26, Dec. 2006.