On routing controls in ISP topologies: A structural perspective

Ryota Fukumoto, Shin'ichi Arakawa, and Masayuki Murata
Graduate School of Information Science and Technology, Osaka University
Email: r-fukumoto@ist.osaka-u.ac.jp, arakawa@ist.osaka-u.ac.jp, murata@ist.osaka-u.ac.jp

***Abstract*— Recent measurement studies on the Internet topology show that connectivity of nodes exhibits power-law attribute, but it is apparent that only the degree distribution does not determine the network structure, and especially true when we study the network-related control like routing control. In this paper, we evaluate the optimal routing and several routing methods on the ISP router level topologies where degree distributions follow power-law. We then examine how structural characteristics of topologies affect the network performance. The evaluation results show that the optimal routing method in a topology obtained by a modeling method considerably increases network throughput. However, the optimal routing method in ISP topologies cannot achieve high network throughput as observed in the modeling-based topology. Our results also show that our proposed routing method achieves almost the same network throughput to the optimal routing method in power-law networks, and more importantly, it exhibits the similar distribution of link utilization.**

I. Introduction

Recent measurement studies on Internet topology show that connectivity of nodes exhibits the power–law attribute [1]. That is, the probability $p(k)$ that a node is connected to k other nodes follows $p(k) \sim k^{-\gamma}$. In recent years, a considerable number of studies have investigated power–law networks whose degree distributions follow the power–law [2]–[6]. A theoretical examination of the characteristics of the power–law network is presented in Ref. [7]. It also presents a Barabashi–Albert (BA) model in which the topology grows incrementally and links are placed based on the connectivity of the topologies to form power–law networks.

Some studies on power–law networks certainly evaluate the distributions of link / node load [5], [8]. Here, the load is defined as the number of node–pairs on those nodes / links. In these studies, the distribution of node load in the topology generated by BA model also follows power–law attribute. However, these studies concentrate on the minimum hop routing. Actually, various approaches have been proposed in order to avoid congestions (i.e., reduce the maximum link utilization) in operating networks. One approach is to introduce MPLS (Multi–Protocol Label Switching) in the networks. Another approach is to optimize link metrics for OSPF (Open Shortest Path First) [9].

In this paper, we evaluate the minimum hop routing and the optimal routing on the ISP router level topologies, where degree distributions follow power–law. We then examine how structural characteristics of topologies affect the routing mechanisms. In particular, we examine the optimal routing and evaluate the distribution of link utilizations from the view point of structual properties. In addition, since the optimal routing usually takes huge time to obtain the solution, we also examine several heuristic routing methods on the topologies. For this purpose, we propose a heuristic routing method that achieve near–optimal throughput in power–law networks. Our heuristic routing method is based on the technology constraints of IP routers [10]. The evaluation results show that our routing method reduces the maximum link utilization about 1/3 in the ISP topology. More importantly, the results show that the distribution of the link utilizations with our routing method is similar to that of the optimal routing method.

This paper is organized as follows. In section II, we propose near–optimal heuristic routing method based on the technology constraints of IP routers. In section III, we evaluate the minimum hop routing, optimal routing, and our proposed routing on router–level ISP topologies. Section IV concludes the paper.

II. A Heuristic Routing Algorithm Suitable to Router–level topologies

Owing to the technology constraint of routers, the degree of nodes and the capacity of links that are connected to the node are highly correlated. That is, when the degree of the node is small, a link that is connected to the node can have a large capacity. On the other hand, when the degree of the node is large, the capacity of corresponding links has to be reduced. Our heuristic routing method determines the routes for each node–pair by considering this technology constraint. This is similar to the approaches in Ref. [10], but different in that the authors in Ref. [10] use this fact for constructing the topology, but we use this fact for routing control.

Our method incrementally determines the route of each node–pair. In determining a route for a node–pair, we use following two policies to incorporate the technology constraint of routers. The first policy is to select the route to avoid the higher–degree nodes. The other policy is to select a link whose link capacity is larger. After we obtain the route between two nodes, the remaining costs of all links are updated based on the selected routes. The updated link cost is in turn used for selecting the route for other node–pairs. Details of our algorithm are as follows.

The proposed routing method determines the route from node i to the other nodes (denoted as j in the following steps). That is, for each node i, we perform the following steps:

Step 1: Set the initial costs for Dijkstra's shortest path (SP) algorithm to all links. The cost of links is set proportional to the degree of the destination node of the corresponding link.

Step 2: For each destination node j, repeat the following sub–steps.

Step 2.1: Determine the route (from node i) to node j by calculating minimum cost path by Dijkstra's SP algorithm.

Step 2.2: Increase the cost for links that are used by the selected route at Step 2.1. The amount of the increase is inversely proportional to the actual link capacities (increase by $\alpha \frac{C_{max}}{C_l}$: α is a parameter, C_l is the capacity of the link, and C_{max} is maximum of the capacities in the topology).

At Step 2.2, the costs for links are increased by inversely proportional to the actual link capacities. That is, we increase the cost to some extent if the link capacity is small, so that sub–sequent node–pairs will not use the lower–capacity links.

III. Evaluation of routing methods on router–level topologies

In this section, we evaluate optimal routing and several heuristic routing methods on router–level topologies. We focused on the network throughput and the load distribution in order to evaluate the routing mechanisms.

A. Network models

To clarify the load characteristics in the router–level topology, we use three topologies; the two ISP router–level topologies, and the BA topology. As ISP topologies, we use the Sprint topology and the AT&T topology measured in Refs. [11]. We also use BA topology generated by the BA model. The Sprint topology has 467 nodes and 1292 links. The AT&T topology has 523 nodes and 1304 links. The BA topology is generated such that the numbers of nodes and links of it are the same as that of the Sprint topology. Note that we have confirmed that the connectivity of nodes for these topologies follow the power–law, but not presented here due to space limitation.

In our evaluation, each node–pair generates the same amount of traffic at a unit time in the three topologies. As Li et al. mentioned [10], constraints with router technology limit the degree (i.e., number of ports in the router) and line speed of a port. Thus, it is important for evaluating the link utilization to determine the link capacity appropriately. Since there is no publicity available information of link capacities of these topologies, we allocate the link capacity based on the Cisco 12416 specification [12], supposing that the link capacity dimensioning is optimized for the minimum hop routing. The details are described in Section III-C.

B. Routing method

In our evaluation, we use four routing methods. They are minimum hop routing, optimal routing, invcap routing, and our proposed routing described in section II. The minimum hop routing is the simplest routing and used for comparison. The invcap routing method is recommended by Cisco Systems. The invcap routing method sets the link costs inversely proportional to the capacities of the links and selects the minimum cost path. The optimal routing method selects optimal routes in order to avoid congestions in networks. Here, we assume that the congestion is avoided when the maximum link utilization is lower in this paper. The optimal routing method is based on the flow deviation method [13]. The details are described as follows.

Optimal routing method: To obtain the optimal link load, we use a flow deviation method [13]. The flow deviation method incrementally changes the flow assignment along feasible and descent directions. Given objective function T, the method set w_l as a partial derivative with respect to F_l, where F_l is the amount of traffic that traverses link l. Then, the new flow assignment is solved by using the shortest path algorithm in terms of w_l. By incrementally changing from the old to the new flow assignment, optimal flow assignment is determined. In this paper, we set objective function T to

$$T = \sum_l 1/(C_l - F_l), \qquad (1)$$

where C_l is the capacity of link l and F_l is as defined above.

C. A method for allocating link capacities

In our evaluations, we allocate the capacities of links based on the technology constraints imposed by the Cisco 12416 router, which has 16 line card slots. When a router has 16 or less connected links, all the links can have 10Gbps capacity. If there are more than 16 links connected to the router, the capacity for one or more of the links should be decreased [10].

However, it is difficult to determine which link capacity should be decreased. Therefore, we allocate the capacities of links in a network so that the amount of traffic between a node–pair is maximized with minimum hop routing method, while satisfying the following two technology constraints imposed by routers.

1) The capacity of a link is chosen from a set {100Mbps, 1 Gbps, 2.4Gbps, 4.8Gbps, or 10Gbps};
2) Each router can handle the traffic up to 320Gbps. That is, the total capacity of links connected with the router is 320Gbps or less.

The first constraint corresponds to the link capacity constraint on routers; the set is chosen from the Ethernet technology for 100Mbps and 1Gbps, and optical transmission technology from 2.4Gbps to 10Gbps. The second constraint represents node capacity constraint on routers. Under these constraints, a router accommodates several low speed tributaries, i.e., it has more than 16 out–going links in the current case, unless the total capacity of the links violates the second constraint.

The algorithm for allocating link capacities is as follows. We give an amount of traffic between node i and node j, d_{ij}, as input values. We assume that each node–pair generates

TABLE I

THE NETWORK THROUGHPUT [GBPS]

	Sprint	AT&T	BA
Minimum Hop	256.85	117.68	364.26
Invcap	114.42	138.60	1157.69
Proposed	405.82	248.95	2444.75
Optimal	627.65	337.34	2706.27
Optimal Ratio	2.44	2.87	7.43

the same amount of traffic at a unit time in the above–mentioned three topologies, that is, we set the identical value d to d_{ij}. Then, we check whether the node capacity constraint is satisfied. If the constraint is violated, we decrease the capacity of links such that the link capacity constraint is satisfied. If there are no allocations of link capacities that satisfy two constraints, d is decreased.

The specific procedure is as follows. Given a network topology, flow assignment (in this case, minimum hop routing paths), and the amount of traffic generated between nodes, d, do following steps;

- Step 1: For all link l, calculate the amount of flow, F_l, that traverses link l by using the given flow assignment and the (identical) traffic demand between nodes d.
- Step 2: Initialize the all link capacity; set the capacity, C_l, of the link l to be 10Gbps.
- Step 3: Check whether the node capacity constraint is satisfied. For each node p, repeat the following steps.
  - Step 3.1: Calculate the sum of the capacity of all links (denoted by C_p^{all}) that are connected to node p.
  - Step 3.2: If C_p^{all} is greater than 320Gbps, i.e., the capacity constraint of node is not satisfied, do the following steps until $C_p^{all} \leq$ 320Gbps.
    - Step 3.2.1: Decrease the capacity of all links that are connected to node p by one step lower than the current capacity, e.g., 10Gbps $\rightarrow$ 4.8Gbps if the current link capacity is 10Gbps. However, we do not decrease the capacity of link if C_l becomes less than F_l.
    - Step 3.2.2: If we cannot decrease the capacity of links any further, decrease the amount of flow between nodes i and j, d_{ij} (in the current case d), and go back to Step 0.

After this algorithm finishes, the capacity C_l for each link l is obtained. We also obtain d_{ij} (=d) as the maximum traffic demand for node i and j at which the network can accommodate.

D. Comparison of the network throughput

In the following section, we evaluate routing methods on the network model described in the previous sections. We first focus on the network throughput. Table I summarizes the network throughput. Here, we define the network throughput as the amount of traffic that the network can accommodate. If the amount of traffic increases more than that value, the maximum link utilization beyond 1.0. In the table, the "Optimal Ratio" represents the ratio of the result of the optimal routing method to that of the minimum hop routing method. The parameters α of our proposed routing method are set optimally for each topology. The values of α are 6, 1, and 1 respectively for the Sprint, the AT&T, and the BA topologies. We discuss effects of this parameter settings later (in section III-F).

Let us first look at the results when minimum hop routing is applied. We observe that the ISP topologies, the Sprint topology and the AT&T topology, have lower network throughput than the BA topology. If the Invcap routing method is applied, the BA topology can accommodate more traffic than the other topologies. Notably, the Sprint topology with Invcap routing method accommodates less traffic than that with minimum hop routing method. Thus, the Invcap routing doesn't work as discussed in [9]. The optimal routing method certainly increases the network throughput. However, comparing with the results by minimum hop routing method, optimal routing in the BA topology (Optimal ratio is 7.43) is more effective than that in the Sprint topology (2.44). The reason is explained as follows. If minimum hop routing is used, the bottleneck link of the BA topology is the link that is connected to the higher–degree node where the link capacity tends to be small according to the technology constraint of routers. Furthermore, lower–degree nodes, which can have the larger capacity for neighboring links, is not fully utilized in the BA topology. If optimal routing is used, lower–degree nodes assist to accommodate traffic. Thus, the maximum traffic demand increases greatly in the BA topology. Our heuristic routing method shows the similar result to the optimal routing. That is, the proposed routing method is also more effective in the BA topology.

E. Comparison of load distribution

We next focus on the load distribution on the router-level topologies from the view point of structural properties. Particularly, we focus on the distribution of link utilizations and the distribution of node load.

1) Link utilization: Figure 1 shows the distributions of link utilization. The vertical axis represents link utilization that is defined as F_i/C_i. The horizontal axis represents the rank of link utilization. The amount of traffic in each topology is obtained by setting d such that the maximum link utilization becomes 1.0 by using the minimum hop routing method. Note that the line captioned "Non degree" and "Non link" are used to discuss variants of our heuristic routing method and not used in this subsection.

First, we observe that, with minimum hop routing, few links are congested while most links are not congested in all topologies. In particular, the AT&T and BA topologies have some highly congested links. However, if we use optimal routing, there are fewer variations in link utilization, i.e., most of links have almost the same utilization. For the Sprint and AT&T topologies, the variations are larger, which is different tendency observed in the BA topology. The reason for this comes from the structure of the Sprint and AT&T topology, where the cluster coefficient is much larger than the BA topology, as shown in Ref. [8]. In other words, a node that is connected to another node (say node A) is also connect to a neighbor (or near, as in terms of physical distance) node from node A. In this case, the congestion at some links cannot be avoided even if another link is selected, because that link is still connected to a node around the congested link. Preferential attachment in the BA model, on the other hand, does not incorporate the locality of connecting nodes, and thus optimal routing can find low–congested links. These results indicate that both the ISP topologies and the BA topology do not fully utilize the link capacity by the minimum hop routing. These results also indicate that the ISP topologies have the structures that give lower throughput working with the optimal routing than the BA topology. As for the proposed routing method, Figure 1 clearly shows that our proposed routing method has the similar distribution of link utilization to the one obtained by optimal routing method.

2) Node load: We next show the distributions of amount of traffic that passes through nodes in Fig. 2. The vertical axis represents amount of traffic that passes through nodes, and the horizontal axis represents the rank of node traffic. Again, the amount of traffic in each topology is obtained by setting d such that the maximum link utilization becomes 1.0 by using the minimum hop routing method.

As in the case of link utilization, we can see that a few nodes are congested while most nodes are not congested with minimum hop routing in these three topologies. Especially, different from ISP topologies, the BA topology has some highly congested nodes; the amount of node traffic is doubled. On the other hand, the variations by the optimal routing are relaxed and the maximum node traffic is decreased. The proposed routing method also shows the same tendency to the optimal routing method.

F. The impact of parameter settings

Our proposed routing method has two heuristics; to select the route to avoid the nodes with higher–degree, and to select links whose capacities are larger. We also use a parameter α to control the extent of these two heuristics. In this subsection, we evaluate how these heuristics affect the network throughput and distribution of link utilization.

We can see the impact of the two policies of our heuristic routing method in Figs. 1 and 2 by looking at the results of "Non degree" and "Non link". "Non degree" means the heuristic routing method does not consider to avoid higher–degree nodes. This routing method is realized by setting the initial costs 1.0 in Step 1 of our heuristic routing method. "Non link" means the heuristic routing method does not take care of link capacities. This routing method is realized by replacing $\alpha\frac{C_{max}}{C_l}$ to α in Step 2.2 of our heuristic routing method.

As for the distribution of link utilizations, "Non degree" is similar to our heuristic routing method in Sprint and AT&T topologies from Figs. 1(a) and 1(b). On the other hand, in the BA topology (Fig. 1(c)), we can see that "Non degree" is much worse than our proposed routing method. We also observe that "Non link" is much worse than our proposed routing method in any topology.

The similar observation can be found in the distribution of node traffic; results of "Non degree" and "Non link" are worse than our proposed routing method in any topology. We therefore conclude that we should control the routes by considering not only capacities of links but also the degree of nodes.

We next show the impact of parameter settings. In Fig. 3, we show the maximum link utilizations. We also show the results when our proposed method is applied on other ISP topologies; Verio topology, and Level3 topology in Ref. [11]. Verio and Level3 are the major ISP providers in United States. Figure 3 clearly shows that the maximum link utilizations is decreased when the parameter α is between 1.0 and 10.0. However in BA topology and Level3 topology, as the parameter α increases, the link utilization increases. This is because the hub–nodes that have many links are located at "the center" of BA topology and Level3 topology, and the traffic therefore tend to concentrate on the hub–nodes. Here, "the center nodes" means a place where other nodes reach the nodes within a few hop count. When the parameter α increases, costs of high bandwidth links also increase. Then the routing method becomes to select shorter hop routes rather than the routes with high bandwidth links. That is, the routing method becomes to select the hub–nodes. However, the links that are connected to the hub–nodes have low bandwidths due to constraints imposed by routers. Therefore, the utilizations of these links increase. To understand these observations more clearly, we show the average hop counts from the highest degree node to the other nodes in Table II. BA topology and Level3 topologies have smaller average hop counts than the other topologies. Therefore, BA and Level3 topologies have structures that higher–degree nodes are located at "the center" of the topology. On the other hand, Sprint, AT&T, and Verio topologies have structures that the higher–degree nodes do not located at the center of the topology, and thus the traffic that passes through the nodes is not much. We therefore conclude that we should set α to 1 or 2. Otherwise the traffic can concentrate on particular nodes.

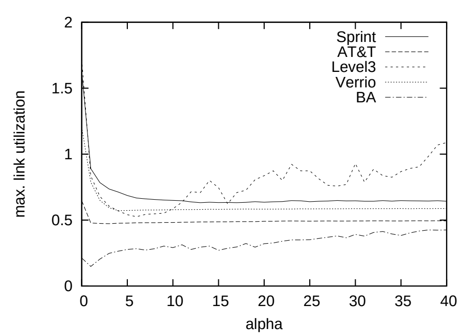

Fig. 3. Effect of α

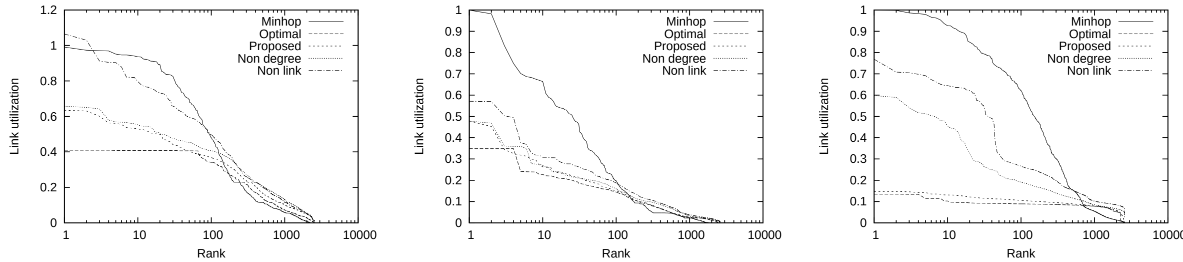

(a) Sprint topology (b) AT&T topology (c) BA topology

Fig. 1. Distributions of link utilization

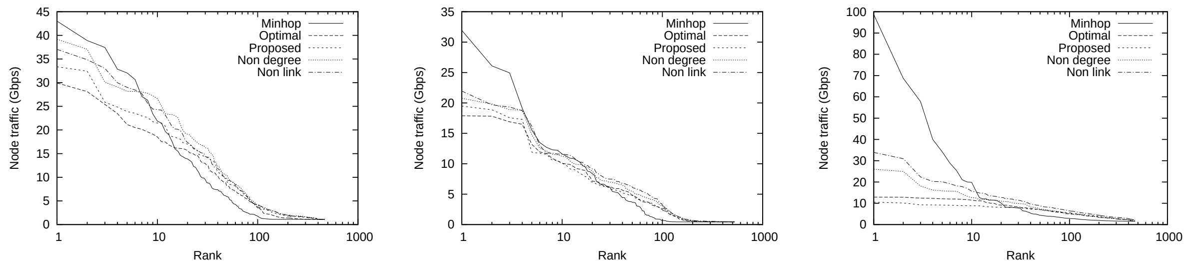

(a) Sprint topology (b) AT&T topology (c) BA topology

Fig. 2. Distributions of node load

TABLE II
AVERAGE HOP COUNT FROM MAXIMUM DEGREE NODE TO THE OTHER NODES

Sprint	AT&T	Verio	Level3	BA
2.89	3.99	3.50	2.22	2.15

IV. CONCLUSION AND FUTURE WORKS

In this paper, we have evaluated several heuristic methods and an optimal routing method in ISP topologies. At first, we have proposed a heuristic routing method with consideration of technology constraints of IP routers, which achieves near–optimal throughput performance and load distribution. The proposed routing method routes between two nodes based on two heuristics; select the route to avoid nodes having higher-degree, and select links that have larger link capacities. We have then evaluated a minimum hop routing method, an optimal routing method and our proposed routing method on ISP topologies, and revealed that how the structural properties of ISP topologies affect the network performance. Since ISP topologies have higher cluster-coefficient and locally connected, the network throughput obtained by optimal routing method was not significant as observed in modeling-based topologies.

In this paper, we have assumed that each node–pair has the same amount of traffic in our evaluations. It is one of future works that we evaluate routing methods and link properties with more realistic traffic demand. Furthermore, the topologies used in computer simulations are the ones measured by Rocket-fuel tools. If we discuss scalability of routing control, we may require a modeling method for generating large-scaled and more realistic ISP topologies, but it is left for our future research work.

REFERENCES

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

[2] A. Akella, S. Chawla, A. Kannan, and S. Seshan, “Scaling properties of the Internet graph,” in *Proceedings of the Twenty-second Annual Symposium on Principles of Distributed Computing*, pp. 337–346, 2003.

[3] W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker, “Scaling phenomena in the Internet: Critically examining criticality,” *Self–organized Complexity in the Physical, Biological, and Social Sciences*, Mar. 2001.

[4] C. Gkantsidis, M. Mihail, and A. Saberi, “Conductance and congestion in power law graphs,” in *Proceedings of ACM SIGMETRICS*, pp. 148–159, June 2003.

[5] K. L. Goh, B. Kahng, and D. Kim, “Universal behavior of load distribution in scale–free networks,” *Physical Review Letters*, vol. 87, Dec. 2001.

[6] L. Qiu, Y. R. Yang, Y. Zhang, and S. Shenker, “On selfish routing in Internet–like environments,” in *Proceecdings of ACM SIGCOMM 2003*, pp. 151–162, Aug. 2003.

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

[8] Shin’ichi Arakawa, Ryota Fukumoto, Tetsuya Takine, Masayuki Murata, “Analyzing and modeling router-level internet topology,” *IEICE Tech. Rep.*(IN2005-95), pp. 43–48, Oct. 2005.

[9] B. Fortz and M. Thorup, “Internet traffic engineering by optimizing OSPF weights,” in *Proceedings of INFOCOM (2)*, pp. 519–528, 2000.

[10] L. Li, D. Alderson, W. Willinger, and J. Doyle, “A first–principles approach to understanding the Internet’s router–level topology,” *ACM SIGCOMM Computer Communication Review*, vol. 34, pp. 3–14, Oct. 2004.

[11] N. Sprint, R. Mahajan, D. Wetherall, and T. Anderson, “Measuring ISP topologies with rocketfuel,” *IEEE/ACM Transactions on Networking*, vol. 12, pp. 2–16, Feb. 2004.

[12] “Cisco systems,” `http://www.cisco.com/`.

[13] L. Fratta, M. Gerla, and L. Kleinrock, “The flow deviation method: An approach to store-and-forward communication network design,” *Networks*, vol. 3, pp. 97–133, 1973.