

Research Article

Analyzing the Evolution and the Future of the Internet Topology Focusing on Flow Hierarchy

Yu Nakata, Shin'ichi Arakawa, and Masayuki Murata

Graduate School of Information, Science and Technology, Osaka University, 1-5, Yamadaoka, Suita, Osaka 565-0871, Japan

Correspondence should be addressed to Yu Nakata; y-nakata@outlook.com

Received 22 December 2014; Revised 5 May 2015; Accepted 12 May 2015

Academic Editor: Lixin Gao

Copyright © 2015 Yu Nakata et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

In the Internet, Autonomous Systems (ASes) exchange traffic through interconnected links. As traffic demand increases, more traffic becomes concentrated on such links. The traffic concentrations depend heavily on the global structure of the Internet topology. Therefore, a topological evolution considering the global structure is necessary to continually accommodate future traffic amount. In this paper, we first develop a method to identify the hierarchical nature of traffic aggregation on the Internet topology and use this method to discuss the long-term changes in traffic flow. Our basic approach is to extract the “flow hierarchy,” which is a hierarchical structure associated with traffic aggregation. Our results show that the current connection policy will lead to a severe traffic concentration in the future. We then examine a new evolution process that attempts to reduce this traffic concentration. Our proposed evolution process increases the number of links in the deeper level in the hierarchy, thus relaxing the traffic concentration. We apply our evolution process to the Internet topology in 2000 and evolve this scenario over 13 years. The results show that our evolution process could reduce the traffic concentration by more than half compared with that without our evolution process.

1. Introduction

The Internet is the largest network system in the world and is becoming ever larger. The amount of traffic on the Internet has been increasing owing to the increase in the number of network users, network services, and communication devices, such as PCs, smartphones, and tablet devices. An AS is a network that is managed by an organization under a single administrative control. The Internet consists of many ASes and the connections among ASes. According to Border Gateway Protocol (BGP) data in [1, 2], the number of ASes has doubled over the last decade; as of November 15, 2013, there were at least 45,980 ASes and 105,540 interconnected links. The number of ASes is estimated to continue increasing in response to the increase in mobile traffic, which doubles each year, and traffic from new and emerging applications using, for example, sensor devices with a communication function [3, 4].

As the traffic amount increases, more traffic will concentrate on existing links. To relax the traffic concentration, each AS tries to form new links with ASes that have not yet been connected. An AS usually has its own policy for

selecting which ASes to connect with from among the many candidates. For example, an AS attempts to connect with another AS such that the cost, revenue, and performance after connecting are optimized. That is, new links are constructed based on the local decision of two ASes. They do not consider the global structure of the Internet topology. However, because the degree of traffic concentration on links depends heavily on the global structure of the topology [5], local decision-making is inadequate to fundamentally avoid the future traffic concentration associated with the increase in traffic. An evolution that considers the global structure of the Internet topology is necessary to continually accommodate future traffic amount.

The evolution of the Internet topology has been studied intensively in recent years [6–8]. Dhamdhare and Dovrolis [6] investigated the long-term change in the number of peering/transit links. The authors also discussed the factors behind the emergence of the current topological structure and gave graph generation models for the Internet topology. Shavitt and Weinsberg [8] used the clustering coefficient [9] and betweenness centrality [10] to characterize this evolution, while Gregoria et al. [7] extracted well-connected

subgraphs from the Internet topology and discussed how these subgraphs were connected to the rest of the Internet topology. These studies longitudinally analyzed the change in the Internet topology from a graph metrics perspective. However, a more important metric to avoid future traffic concentration is related to the change in the structure of the Internet topology associated with spatial dynamics of the traffic flow. An analysis of the change in the structure associated with traffic flow can help to reveal where the traffic concentration occurs and how to deal with it.

We therefore develop a method to identify the hierarchical nature of traffic aggregation in the Internet topology and use this method to discuss the long-term changes in traffic flow. Our basic approach is to extract the “flow hierarchy,” which is a hierarchical structure associated with traffic aggregation, from the Internet topology. Many works have shown that the Internet has a hierarchical structure [11–13]. Within this hierarchical structure, an AS aggregates traffic from lower-level ASes and relays the traffic to higher-level ASes. Such traffic aggregation leads to a hierarchy of traffic aggregation, which in turn leads to the traffic concentration on links. Recently, the structure of the Internet topology is becoming “flat” [14], and the trend of traffic flow also is changing from centralized to more distributed. Nevertheless, the flow hierarchy has not disappeared because the flat structure is formed by adding links to existing hierarchical structure. To extract the flow hierarchy, we focus on structures called “modules” as the unit of traffic aggregation and retrieve the hierarchy of modules that appear in the Internet topology. A module consists of a set of ASes that are densely connected with each other, and each module is sparsely connected with other modules [15]. The outgoing traffic from one module is first aggregated inside that module, and then the traffic is transferred to the other module through the sparsely connected links. A module may be divided into two or more submodules; that is, there is a containment relationship between the module and submodules. By repeating the division of modules and revealing their containment relationships, we can extract the flow hierarchy of the Internet topology. We then investigate the long-term changes in the flow hierarchy of the Internet topology. Our results show that the increase in traffic amount at the top-level module is larger than that at middle-level or low-level module and particularly has slightly accelerated since 2011. This suggests that the current connection policy will lead to a severe traffic concentration in the future Internet topology. Therefore, we urgently need an evolution process that considers the global structure of the Internet topology to slow down the increase in traffic concentration. In this paper, we examine a new evolution process that attempts to increase the number of links between lower-level modules to relax the traffic concentration in higher-level modules. We apply our evolution process to the Internet topology in 2000 and evolve this scenario for 13 years. We then evaluate the traffic concentration at various levels of containment following the evolution. The results show that our evolution process can suppress the traffic concentration by more than half compared with that without our evolution process.

This paper is organized as follows. Section 2 gives an overview of some related work in the analysis of the Internet topology. Section 3 describes the hierarchy concept based on the containment relationship of modules and presents the method of extracting the flow hierarchy from the Internet topology. Section 4 discusses the long-term change in the flow hierarchy of the Internet topology. We first investigate the internal structure in a module and then illustrate the structure between top-level modules in the flow hierarchy, because a large amount of traffic traverses the links between top-level modules. Finally, we investigate the long-term change in the structure of each level in the flow hierarchy. Section 5 studies the links on which a lot of traffic is aggregated. In Section 6, we examine a new evolution process that attempts to increase the number of links between lower-level modules. We apply the evolution process to the Internet topology in 2000 and confirm that it suppresses the traffic concentration across links between top-level modules. Section 7 shows that the appearance of Hyper Giants does not enable the continued accommodation of an increase in traffic amount. Section 8 concludes this paper.

2. Related Work

Understanding and analyzing the structure of the Internet topology is important, because the properties of the Internet are used for network design. The network performance, such as the amount of traffic that can be accommodated across the Internet, is dependent on the structure of the Internet topology, because this strongly affects the traffic flow. Therefore, when a network operator of an AS adds new links and network equipment, a design based on the properties of the topology is needed to improve the network performance. Determining the structure of the Internet topology is also vital to evaluate the performance of new applications and protocols on a topology reflecting the structure and properties of the Internet. For example, a topology reflecting properties of the Internet is required to evaluate the scalability of BGP [16].

For the past dozen years, various structural properties of the Internet topology have been widely investigated. References [17, 18] visualized the Internet topology to determine its structural properties. However, it is difficult to characterize structural properties from pictures of the Internet topology generated by these studies, because the Internet topology is large and complex. Some studies have investigated structural properties using various graph metrics. In [19], Faloutsos et al. revealed that the degree distribution of the Internet topology exhibits power-law attributes, and Pastor-Satorras et al. [12] showed that the distribution of betweenness centrality also follows a power law. However, these studies analyzed the structural properties at a point in time. Network design requires the prediction of the future structure of the Internet topology. To predict the future structure, the trend in changes to the Internet topology has to be clarified. In [6], Dhamdhere and Dovrolis quantified the ability of an AS to attract customer ASes that pay a transit fee for traversing traffic and found that Internet Service Providers (ISPs) connecting to a lot of customer ASes had acquired more customer

ASes. These studies analyzed the evolution of the Internet topology using some graph metrics. Each graph metric shows a characteristic of the Internet topology; however, these are not directly related to network performance. For instance, even if two networks have the same degree of distribution, the amount of network equipment needed to accommodate traffic demand will differ depending on the structure of the networks. For example, [5] found that the degree of traffic concentration on links is heavily dependent on the global structure of the topology. Actually, the Internet topology suffers from traffic congestion more than a random network [20]. It is important to understand the global structure related to the spatial dynamics of traffic flow to develop a new evolution process that avoids the current and future traffic concentration suffered by the Internet topology.

In [8], Shavitt and Weinsberg analyzed changes in topological structure, such as betweenness centrality and link density, by focusing on large content providers, also referred to as Hyper Giants [14, 21]. From this analysis, it was found that the structure of the Internet topology has changed from a hierarchical to a flat structure. This is because large content providers construct links with a lot of small ISPs. Because they have influenced the Internet topology, considerable attention is currently focused on these Hyper Giants. However, Hyper Giants do not contribute to the moderation of traffic concentration over certain parts of the links, because the traffic flow between two ASes does not traverse the Hyper Giants; that is, the traffic is not aggregated at the links controlled by the Hyper Giants. Thus, the Hyper Giants are not relevant to an evolution process to reduce traffic concentration at these links. In this study, we focus on the structure of traditional links, such as those between ISPs.

3. The Flow Hierarchy

3.1. Concept of the Flow Hierarchy. We use the flow hierarchy to reveal where and how traffic is aggregated. The structure of the flow hierarchy is the hierarchical structure based on containment relationship of modules. We note that the flow hierarchy is not a hierarchy of “tier” based on the ISP’s business scale but the structure indicating a gradual traffic aggregation in the Internet topology. Such a containment relationship has appeared in the history of the Internet evolution, and then the traffic is aggregated in accordance with the flow hierarchy. This makes the flow hierarchy be useful for analyzing degree of traffic aggregation. In the late 1960s, some academic organizations deployed network equipment and connected with all other organizations. This is the origin of the Internet, and the organizations became to be called ASes later. To participate in the early Internet, new ASes needed to connect with all other ASes. However, as the scale of the Internet became larger, it was increasingly difficult to sustain the full mesh network. Because the construction and maintenance costs of long or high-capacity links are high, new ASes tend to connect with only a few “senior” ASes that have long or high-capacity links. As a result, sets of ASes centered on senior ASes, that is, modules, were generated. However, as the number of ASes connecting to senior ASes increases, the amount of traffic aggregated at senior ASes

and global links increases, and the risk of suffering traffic congestion increases. To reduce the traffic load at senior ASes, some ASes have locally aggregated traffic. Because a hierarchical structure has appeared in the Internet under this process of traffic aggregation, the flow hierarchy reflects the hierarchical nature of traffic aggregation. Therefore, we use the flow hierarchy to reveal where and how traffic is aggregated.

3.2. Extraction of the Flow Hierarchy. We now extract and investigate the flow hierarchy in the Internet topology. First, we obtain the topology data of ASes and links in the Internet topology (Section 3.2.1). We then extract the hierarchical structure based on containment relationship of modules from the Internet topology (Section 3.2.2). We finally give the traffic demand to the hierarchical structure because the flow hierarchy is derived by adding the traffic amount on each link to the hierarchical structure (Section 3.2.3).

3.2.1. Obtaining Topology Data. We obtain the topology data of ASes and links in the Internet topology. We extract the topology data from the BGP routing tables that have been recorded in the gateway routers of large ISPs and have been gathered. Various organizations, such as UCLA [22] and CAIDA [17], create the Internet topology data and these topology data include more links [23, 24]. However, these topology data are not suitable for a longitudinal topological analysis because the number of monitors that observe BGP tables and traceroute results that are used to create topology data has greatly increased. That is, we cannot distinguish between actual evolution of the topology which is contributed by the real change of the topology and changes caused by the increase of monitors. Instead of the data provided by UCLA and CAIDA, we use the BGP tables gathered by a part of servers of RouteViews Project and RIPE NCC. The part of servers has been gathering the BGP tables from almost the same ISPs after starting up their projects. Although the number of links observed is fewer than the topology data of UCLA and CAIDA, the BGP tables from RouteViews Project and RIPE NCC are suitable for a longitudinal topological analysis because they are consistently comparable over time. BGP tables contain AS paths, which are the routes between two ASes. The AS path is described as a list of ASes that the traffic traverses. From the AS paths in the BGP tables, we obtain the ASes and links in the Internet topology. We use BGP routing tables stored in <http://archive.routeviews.org/oix-route-views/>, which is a RouteViews Project server, and rrc00.ripe.net, which is a RIPE NCC server. The reason why we use these servers is that they are the oldest ones that are still working. Table 1 shows the number of ASes and links that we can extract. Unfortunately, [6, 25] reported that this method cannot capture over 40% of the peering links on which traffic is exchanged without a transit fee. Since a huge amount of traffic traverses peering links through IX, missing of peering links decreases the accuracy of estimated amount of traffic traversing each link. However, the use of BGP routing tables is not a problem for this study because the purpose of this study is to reveal the impact of the global structure of

TABLE 1: Number of ASes, links, and AS paths in the Internet topology.

	June 15, 2000	June 15, 2004	June 15, 2008	November 15, 2013
# of ASes	8,162	18,015	29,320	45,980
# of links	17,533	40,205	64,305	105,540
# of AS paths	299,434	1,108,704	1,901,745	3,136,820

the Internet on the traffic concentration rather than to show the actual traffic amount.

3.2.2. Extracting the Hierarchy of Modules Based on Containment Relationship. The structure of the flow hierarchy is the hierarchical structure based on containment relationship of modules. We extract the hierarchical structure from the Internet topology. The hierarchical structure based on containment relationship of modules is extracted by repeating the division of modules into submodules. Several methods for division of modules have been proposed such as the Infomap method [26], the OSLOM method [27], and the Louvain method [28]. Since our main concern is the traffic aggregation, we select the Louvain method for our analysis. The Infomap method uses the probability flow of random walks on a network as a proxy for information flows in the real system and divides the network into modules by compressing a description of the probability flow [26, 29]. However, since the traffic flow of the Internet is not random walk, we cannot capture links where the traffic is aggregated from the Infomap method. The OSLOM method uses a measure indicating how obvious module structure is in the network against a random null model graph. Therefore, the OSLOM method can detect the obvious module structure against the random null model graph. However, the traffic concentration is expected to be observed in also the random null model graph and the OSLOM method cannot capture the traffic concentration. Unlike the Infomap and the OSLOM methods, the Louvain method derives modules so that the number of intermodule links relative to that of the intramodule links is minimized. The traffic originated inside a module is first conveyed and aggregated by the intramodule links and then transferred by the intramodule links. The Louvain method incrementally merges modules into a module, so we can gradually capture links where the traffic is aggregated by using the Louvain method.

In the Louvain method, the topology is divided in such a way as to maximize the modularity. The modularity is a measure of the strength of interconnection among modules when a particular division P of a topology is given and is defined by

$$M(P) = \frac{1}{2m} \sum_{ij} \left[A_{ij} - \frac{k_i k_j}{2m} \right] \delta_{S_i, S_j}. \quad (1)$$

Descriptions of the variables in (1) are shown in Table 2. Here, we regard the maximum of $M(P)$ for all divisions as

TABLE 2: Definition of notation for modularity (1).

Notation	Definition
m	Number of links
i, j	Node
A_{ij}	Element of adjacency matrix
k_i	Degree of node i
S_i	Module that contains node i
δ_{S_i, S_j}	Kronecker delta. If S_i and S_j are the same, δ_{S_i, S_j} is 1, otherwise 0

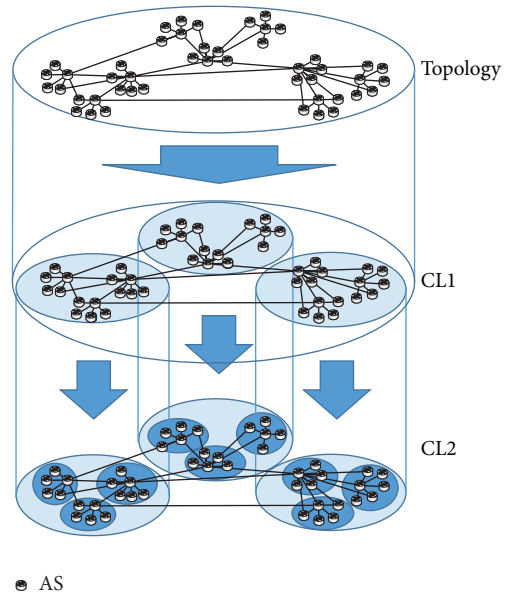


FIGURE 1: Extraction of hierarchical structure based on the containment relationship of modules.

the modularity of the topology. The modularity of a topology ranges from 0 to 1. The modularity is high in case that links between ASes in the same module are densely connected and links between ASes in different modules are sparsely connected. The modularity of a complete graph and a star graph is 0, because these graphs do not consist of sets of nodes densely connected to each other.

After dividing the Internet topology into modules as described above, we divide each module into smaller submodules. Furthermore, we divide these submodules into even smaller submodules. By repeating this dividing process, the hierarchical structure based on the containment relationship is extracted. If the modularity of a module is 0, it cannot be divided into submodules because the module does not consist of sets of densely connected nodes. All modules are repeatedly divided until their modularity is 0. We define the ‘‘containment level’’ (CL) as the level of the hierarchical structure. As shown in Figure 1, CL1 modules are modules that are extracted in the first division of the Internet topology. Submodules of CL1 modules are CL2 modules, and submodules of CL n modules are CL $n + 1$ modules, where n is a nonnegative integer.

3.2.3. Assigning Traffic Demand. We assign traffic demand to the hierarchical structure of the containment relationship of modules. Since the actual traffic amount on most paths is closed information, we give the traffic demand based on the gravity model [30]. The gravity model is a simple method for estimating the traffic demand [30, 31] and is used in some studies [32, 33]. The traffic demand of AS i is proportional to the degree of AS i since the business scale of AS i is related to its degree [5, 34]. Note that, as discussed in [35], the gravity model does not capture self-similarity and long-range dependence of traffic. However, we use the gravity model for assigning traffic demand since our study focuses on the increase in the degree of traffic concentration rather than short-term traffic fluctuation. The gravity model is represented by the following expression:

$$e_{ij} = \gamma \cdot x_i \cdot x_j, \quad (2)$$

where e_{ij} is traffic amount on the path between AS i and AS j . x_i and x_j are the traffic demand of AS i and AS j , respectively, and γ is a scaling factor and is set to 1 hereafter. Note that this setting may not reflect actual traffic amount. However, our focus here is to reveal the traffic concentration on some links rather than actual traffic amount on each link.

Note that Hyper Giants send huge amounts of traffic compared with the other ASes. In particular, Google and Akamai are defined as Hyper Giants by some studies [5, 14]. We check names of organizations managing ASes in CIDR report [36], and we regard ASes whose names contain ‘‘Google,’’ ‘‘Akamai’’ as Hyper Giants. Then, γ is set to 1 if both of AS i and AS j are not Hyper Giants; otherwise γ is set to 895. These values are determined based on a Cisco report [37, 38] that quantifies the traffic amount on the Internet. Cisco reported the traffic over the whole of the Internet to be 369 exabytes in 2011 and that between users and data centers to be 116 exabytes. The number of ASes registered by the Internet Registry is 60538, and the number of famous content providers is about 30. Therefore, the average amount of traffic at each AS is 4.18 petabytes ($= (369 - 116)/(60538 - 30)$ exabytes), and the average amount of traffic sent by a large content provider is 3.74 exabytes ($= 116/30$ exabytes). Thus, we set γ to 895 ($= 3.74 \cdot 1000/4.18$) for the Hyper Giants. The value of γ may not reflect actual traffic amount. However, our focus here is to reveal the traffic concentration on some links rather than actual traffic amount on each link. Note that Microsoft is also called as Hyper Giants in some studies [5, 14]; however, we regard only Google and Akamai as Hyper Giants because the degree of Microsoft (AS number is 8075) has greatly changed.

When we derive the amount of traffic that traverses each link from the amount of traffic between AS pairs, the path between two ASes is required. Unfortunately, most paths between AS pairs are undocumented. Thus, we assume that the path between two ASes is a minimum hop path, although this is not always the case for actual BGP routings [39]. This assumption is sufficient to observe the change in traffic aggregation, because [8] reported that the number of minimum hop paths that traverse an AS is similar to the actual amount of traffic that traverses that AS. Thus, we

consider minimum hop paths to be useful for analyzing the changes in traffic aggregation.

4. Long-Term Change of the Flow Hierarchy

The traffic concentration at links between modules is dependent on the structure within modules and the structure between modules. An investigation of the change in the flow hierarchy is important for discussing the future evolution of the Internet topology. Therefore, in this section, we first analyze the internal structure of modules. Through this analysis, we investigate ASes that have a lot of links between modules to confirm where traffic is aggregated in the module. We then analyze the between-module structure. In particular, we investigate the structure between CL1 modules, which are top-level modules in the flow hierarchy, because it is thought that large amounts of traffic are aggregated in the links between CL1 modules. Finally, we investigate the long-term change in the structure of each level in the flow hierarchy to reveal the trend of traffic aggregation at links between modules at each level.

4.1. Internal Structure of Modules. In this section, we reveal the internal structure of modules using various graph metrics. Here, we do not divide CL5 modules even when some CL5 modules can be divided to CL6 modules. The reason is that the number of CL6 modules and the size of CL6 modules are too small to see the change of internal structure of the CL5 module. Figure 2 shows the longitudinal change in the graph metrics of modules. From the analyses in Figure 2, we confirm that the internal structure of modules gives a star-like graph. Figure 2(a) shows the mean ratio of ASes that have degree one or two. For the modules in most CLs, the ratio of these ASes is over 80%. Figure 2(b) shows the mean of the maximum degree in the modules. Whereas over 80% of ASes in a module have only one or two links, the degree of hub ASes in CL1 modules is over 100, and that in CL2 modules is over 30. Furthermore, the degree of hub ASes in these CLs has been increasing. Although the degree of hub ASes in CL4 and CL5 is less than 10, they connect to most of the ASes in the module. Thus, it is found that the degree of most ASes in a module is small, but the degree of some ASes is large. To reveal where the links are constructed in a module, we now investigate the assortativity of modules. Assortativity is an index indicating that a node in a network connects with ASes that have a similar degree [40]. If all links are constructed between nodes that have the same degree, the assortativity of the network is 1. On the contrary, the assortativity is 0 when there is no correlation between the degrees of two nodes that are connected with each other. When the nodes with small degrees are likely to be connected to the nodes with large degrees, the assortativity is close to -1 . As shown in Figure 2(c), the assortativity of modules in all CLs is small, which means that ASes that have a small degree connect to hub ASes. Figure 2(d) shows the clustering coefficient of modules in each CL to investigate the connection between neighbor ASes. The clustering coefficient of an AS is an index indicating the ratio of connected pairs to all neighbor nodes’

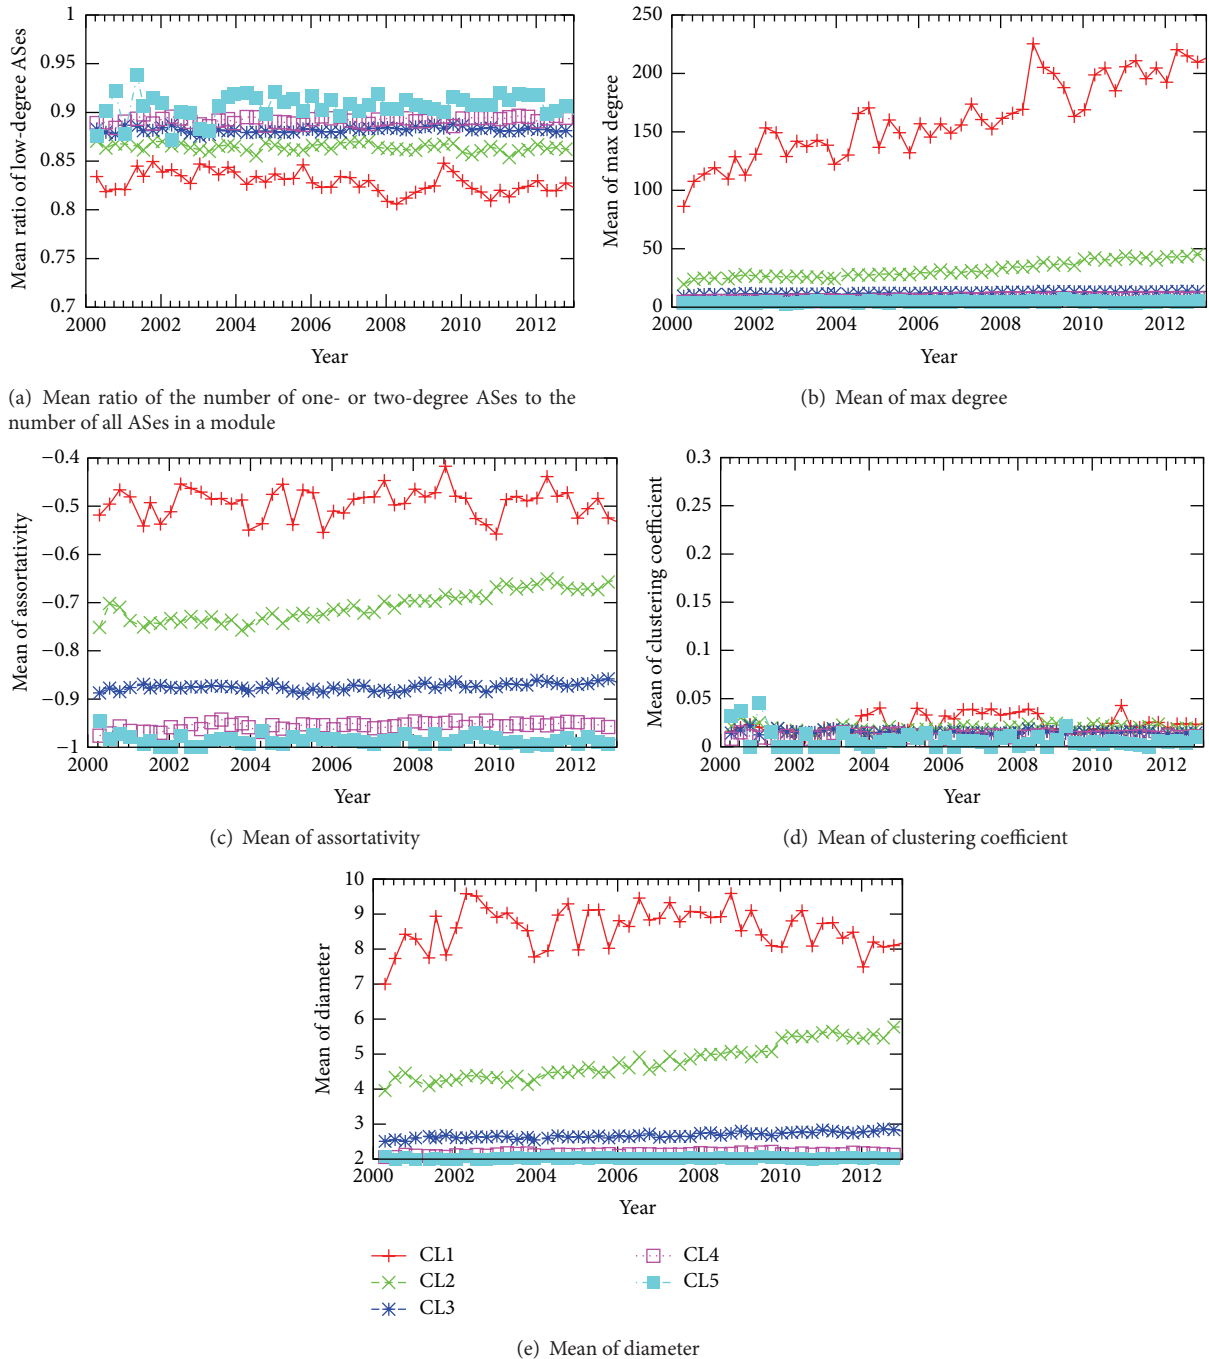


FIGURE 2: Mean value of graph metrics of modules in each CL.

pairs of an AS and ranges from 0 to 1. As shown in Figure 2(d), the clustering coefficient of each module is small. This means that neighbor ASes of a given AS do not connect with each other. The result suggests that hub ASes link to a lot of ASes that have a small degree, and ASes with a small degree are not connected to each other. Finally, Figure 2(e) shows the mean diameter of modules. As the CL increases, the module diameter approaches 2. To summarize the points in Figure 2, it is obvious that the internal structure of each module is a star-like graph.

4.2. *ASes with a Number of Links between Modules.* The large amount of traffic that is generated within a module is aggregated at links between modules. We reveal the relationship between the degree of ASes and the number of links between modules. We first investigate whether a hub AS or an AS that has a small degree has more links between modules. Here, we define hub ASes as those having a degree that is more than half of the maximum in the module. Table 3 shows the average ratio of links between modules to all links of an AS in the Internet topology on 15 July 2013. This shows

TABLE 3: Ratio of intermodule links of hub ASes and ratio of intermodule links of low-degree ASes.

	Hub ASes	Low-degree ASes
CL1	0.244	0.137
CL2	0.396	0.269
CL3	0.476	0.293
CL4	0.463	0.267
CL5	0.406	0.093

TABLE 4: Average number of links between modules.

	e_1	e_2	e_3	e_4	e_5	e_6
Tier-1	989.54	158	624.31	0.54	22.77	0
Sub-Tier-1	191.03	75.23	160.82	2.31	14.31	0
Tier-2	52.76	29.39	52.63	0.93	7.74	0.03
Tier-3	16	11.33	33.08	0.58	2.42	0.17

that hub ASes have a higher ratio than low-degree ASes. This means that low-degree ASes tend to connect to only the hub ASes belonging to the same module. The hub ASes link to both ASes in the same module and ASes in other modules. Therefore, the traffic between modules is aggregated at the hub ASes and then transferred to ASes in other modules.

Next, we examine which types of ASes have many links between modules. In the Internet, there are various types of ASes. ISPs are classified into four types, Tier-1 to Tier-3 and sub-Tier-1. We define sub-Tier-1 as ISPs for which there is no consensus as to whether they should be categorized as Tier-1 or Tier-2. The other ASes are classified as Hyper Giants or Academic. In this study, ASes are ranked based on two types of links: transit links and peering links. A transit link is one in which traffic is exchanged with a transit fee. A peering link is one where traffic is exchanged without a transit fee. Unfortunately, information about the type of link is generally unknown. The method of [17] can infer the link type with an accuracy rate of 99.1%, and so we use this approach. We classify ASes based on the following steps. First, we extract “peering links” and ASes that have peering links. We regard a connected component consisting of peering links as one tier, because two ASes connected with a peering link generally process the same amount of traffic. Next, we check the commercial name of the AS in each connected component and determine the tier of each connected component from six types. Finally, we regard the tier of the connected component that contains the AS to be the type of AS. There is a hierarchy in the Internet based on link types [11–13]. Note that a hierarchy based on AS types is different from the flow hierarchy. The hierarchy based on link types shows the difference in the amount of traffic exchanged by two connected ASes. The flow hierarchy describes the amount of traffic aggregated at ASes or links based on the global structure of the topology.

Table 4 shows the average number of links between modules for each AS type. e_x is the average number of links between modules of an AS in CL x . Table 4 shows that e_x decreases as x increases. This means that Tier-1 ASes




FIGURE 3: Outline drawing of the flow hierarchy in the Internet topology.

have more links between modules than other ASes, because most Tier-1 ISPs have a global network spanning multiple continents and connected with many ISPs all over the world.

According to our findings, the flow hierarchy can be illustrated as in Figure 3. In Figure 3, the number of ASes, number of links between ASes in different tiers, and number of links between ASes in the same tier are 1/5 of those in the actual Internet topology in 2012. In Figure 3, ISPs are arranged from top to bottom in descending order of amount of traversing traffic, and the triangles represent modules. As shown in Figure 3, there is a hierarchy in the Internet based on AS type. Note that this hierarchy is different from the flow hierarchy. The major difference is that the hierarchy based on AS type is not reflected by the structure of the topology. Each module contains ASes in different tiers, and ASes in higher tiers have more links between modules. A module in the flow hierarchy is a part of a vertically divided Internet topology. From the structure in Figure 3, we can see that Tier-1 ISPs exchange traffic traversing from/to other modules. The traffic concentrates at Tier-1 ISPs, because they aggregate the traffic that is generated in the modules.

4.3. Long-Term Change in Structure of Top-Level Modules in the Flow Hierarchy. A hub AS connects links to a lot of ASes in other modules. Thus, a hub AS aggregates traffic generated in the module and relays the traffic to the other modules. Therefore, it is considered that an immense amount of traffic generated in CL1 modules is aggregated at the links between top-level (CL1) modules. In the future, when traffic concentrates at links between CL1 modules, an evolution process is needed that avoids traffic concentration, allowing the Internet topology to accommodate this increase in traffic amount. Thus, the change of traffic concentration at links between CL1 modules must be clarified. For this purpose, we analyze the long-term change in the structure between CL1 modules, because the degree of traffic concentration at the links depends on the connections between CL1 modules.

We first investigate the long-term change in modularity of the Internet topology to investigate the structure between CL1 modules. Since the value of modularity itself is not suitable measure to investigate the modular structure [41–43], we compare the modularity between the ER random model,

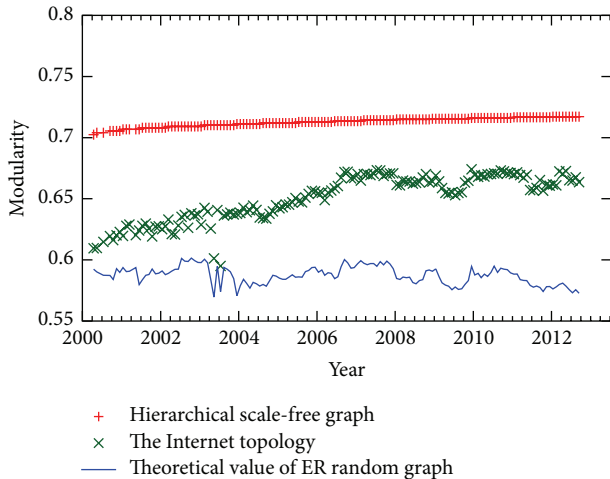


FIGURE 4: Long-term change in modularity.

hierarchical scale-free graph, and the Internet topology. Figure 4 shows the long-term change in modularity of these graphs. For the hierarchical scale-free graph, a module having a scale-free degree distribution is first generated, and this is incrementally added to the graph until the numbers of nodes and links exceed those of the Internet topology. We create the hierarchical scale-free graph that has the same number of nodes and links as the Internet topology. In Figure 4, we use (19) in [41] for calculating the modularity of the ER random model. Equation (19) is an equation to analytically calculate the maximum modularity of the ER random model without the module detection, and the modularity derived by (19) in [41] is close to the modularity derived by the simulated annealing method [41]. There is another approach to calculate the maximum modularity [43]. However, the modularity derived by the equation in [41] is slightly closer to the modularity derived by the simulated annealing method in case that the average degree is fewer than 10 (see Figure 12 in [43]). Since the average degree of the Internet topology is also fewer than 10, we use the equation in [41].

As shown in Figure 4, the hierarchical scale-free graph has the largest modularity, and the ER random graph has the smallest. In the hierarchical scale-free graph, a large amount of traffic tends to be aggregated at the links between modules, because the large modularity indicates a low density of links between modules. In Figure 4, the modularity of the hierarchical scale-free graph and the ER random graph remains constant. On the other hand, the trend in modularity of the Internet topology changed sometime around 2007. This suggests that the overall structure started to change at this time. The dashed line in Figure 4 denotes 1 January 2007. Until this point, the modularity of the Internet topology had been increasing. This suggests that new links had tended to be locally constructed between two ASes in a module. It is thought that when new ASes are created in the Internet topology, they connect to ASes having links between higher-level modules. Since 2007, the modularity of the Internet topology has remained constant. However, the number of links between modules has increased since 2007.

To clarify the factors affecting the change in the modularity trend around 2007, we investigate the long-term change of variables in the definition of modularity (1). The modularity depends on the ratio of links between nodes in a module to all links and the node degree in each module. The key terms in (1) are

$$\alpha = \frac{1}{m} \sum_{ij} A_{ij} \delta_{s_i, s_j},$$

$$\beta = \frac{1}{m} \sum_{ij} \frac{k_i k_j}{2m} \delta_{s_i, s_j}. \quad (3)$$

α is the ratio of links between two nodes in a module to all links, and β is the probability of drawing a link between nodes that are in the same module when the link is randomly deployed on the topology. The higher the degree of node i and node j , the higher the value of β . α and β are normalized by the number of links in the Internet topology. Figure 5 shows the long-term change in these terms. Figure 5(b) shows that β has decreased continuously since 2000. As there is no change in this trend around 2007, β is not considered to be a factor in the change in modularity. On the other hand, Figure 5(a) shows that the trend of α changed around 2007. α was increasing until 2007, with minor fluctuations, and decreases after 2007. Thus, we assert that the change in the trend of α affected the modularity of the Internet topology. Even though the scale of the Internet topology has increased since 2007, the ratio of links between nodes in a module has decreased; that is, the number of links between top-level modules has increased. We believe that the factors behind the increase in links between top-level modules are the reduction in the price of constructing links and the increase in IXes (Internet eXchanges), which are relaying points for traffic between two connected ASes. These factors lead to an increase in intermodule links between ASes that do not have a lot of links between top-level modules. As a result, the modularity of the Internet topology has decreased.

4.4. Long-Term Change of Each Level in the Flow Hierarchy.

More links between modules are needed to avoid an increase in traffic concentration at links between top-level modules. New links between modules should be constructed between two ASes that locally aggregate traffic. This is because a part of the traffic that traverses the existing links between top-level modules will traverse links between two ASes that locally aggregate traffic. We investigate the traffic aggregation at links between modules in each CL to reveal where the ASes that locally aggregate traffic are located. The degree of traffic aggregation at the links between modules in each level of the flow hierarchy depends on the structure between the modules in each level. Therefore, in this section, we investigate the long-term change in the structure of each level in the flow hierarchy.

There are two ways in which the flow hierarchy can evolve: by expanding in depth and by expanding in width. There are two further subcategories for the expansion of the width. One is to increase the number of modules in each CL, and the other is to increase the number of ASes in each module.

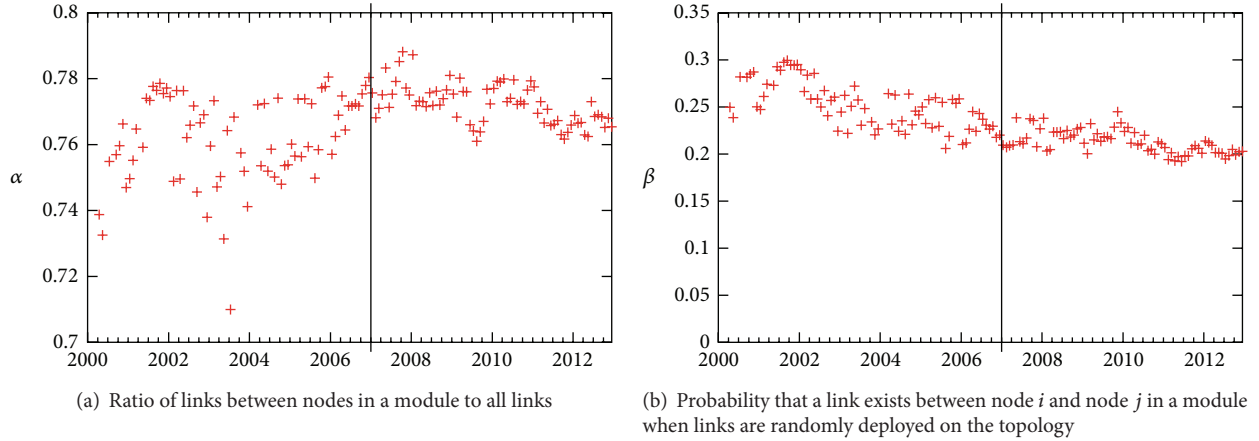


FIGURE 5: Long-term change of variables in the definition of modularity.




FIGURE 6: Evolution of the flow hierarchy with the increase of modules, the number of ASes in a module, or width.

Figure 6 illustrates these expansions of the flow hierarchy. White nodes indicate ASes that exist before the growth, and red nodes indicate those added after the expansion. In the left-hand growth pattern in Figure 6, the number of modules in each CL increases as the topology grows. In this case, the links between modules also increase in number. By increasing the links between modules, the concentration of traffic at existing links is relaxed. In the center growth pattern in Figure 6, a star-like graph in each module becomes larger because additional ASes connect to the hub ASes in each module. As a result, the amount of traffic aggregated at hub ASes and on links between modules increases. In the right-hand growth pattern in Figure 6, submodules are generated in each module. The generation of submodules increases the maximum number of CLs, which corresponds to the depth of the flow hierarchy. If the depth of the flow hierarchy in the Internet topology grows, the amount of traffic aggregated on the links between top-level modules will decrease. This is because the paths between ASes belonging to the same module do not traverse the links between modules in the upper CL.

We first investigate whether the depth of the flow hierarchy has been expanding or not. The depth of the flow hierarchy is defined by the containment level where a module at the level cannot be divided into submodules. Hereafter, we call modules that do not have submodules terminal modules. Figure 7 shows the number of terminal modules at each CL. The value of y -axis is normalized by the total number of terminal modules. We observe that most of terminal modules are located at CL3 and CL4, and the depth of these modules has increased from 2000 to 2012. However, the increase in terminal modules in CL4 is only 10%. Moreover, the average depth of terminal modules is slightly increased, from 3.42 to 3.71. The depth of the deepest terminal module remains steady at six from 2003 to 2013. Therefore, we conclude that the depth of the flow hierarchy has not changed greatly.

We next investigate whether the growth in the flow hierarchy has followed the left-hand pattern or the center pattern in Figure 6. Table 5 shows the number of modules in each CL. The number of modules in CL3 and CL4 is greater than that in other CLs. Furthermore, the number of modules at CL3 and above has increased more rapidly than the number




FIGURE 7: Containment level where a module at the level cannot be divided anymore.

TABLE 5: Number of modules in each CL.

Year	CL1	CL2	CL3	CL4	CL5	CL6
2000	26	255	812	529	71	0
2002	37	384	1215	881	131	6
2004	43	462	1526	1173	208	6
2006	40	479	1883	1562	299	4
2008	40	508	2437	2088	389	6
2010	42	490	2795	2641	438	2
2012	51	578	3153	3181	638	12

TABLE 6: Mean number of ASes contained in one module.

Year	CL1	CL2	CL3	CL4	CL5	CL6
2000	224.97	24.78	8.41	4.01	3.22	2
2002	424.32	35.73	9.81	4.32	3.35	2.75
2004	414.07	37.84	10.41	4.67	3.59	3.5
2006	528.33	43.54	11.25	5.43	3.59	2
2008	695.98	57.18	12.20	6.20	3.71	2
2010	841.98	69.16	12.26	5.71	3.81	2.17
2012	915.93	73.31	12.93	6.36	3.74	2.89

at CL1 and CL2. This means that the structure in CL3 and above has grown in similar fashion to the left-hand pattern in Figure 6. Table 6 shows the average number of ASes in a module. From 2000 to 2012, the average number of ASes in CL1 modules increased by a factor of 4.07, and that in each CL2 module increased 2.96 times. The number of ASes in modules in these CLs increased at a faster rate than in the other CLs. This suggests that the structure in CL1 and CL2 has expanded by increasing the number of ASes within a module. That is, the structure of these CLs has expanded according to the center pattern in Figure 6. The expansion

in width with the increase of ASes in a module leads to an increase in the amount of traffic aggregated at links between modules. Therefore, more traffic has been concentrated at links between CL1 modules and links between CL2 modules. Note that it is known that the Louvain method suffers from a resolution limit. The resolution limit is the characteristic scale of the smallest size of a module that the method can detect. We checked the effect of the resolution limit by comparing divisions by the Infomap method, which is known to mitigate the resolution limit better than the Louvain method [44]. We found that the division by the Louvain method is affected by the resolution limit: the number of small-size (<10 ASes) modules is about ten times fewer than that by the Infomap method. However, we also found that the impact of the resolution limit on analyzing the evolution of flow hierarchy is marginal (see Appendix A for detail). The main reason is that the evolution of the flow hierarchy depends on the relation between the large-size module at the CL and the large-size module at lower-level CL. That is, the evolution of flow hierarchy indicates how the large-size module at a CL can be divided into submodules at lower-level CL. Our result shows that the resolution limit of the Louvain method is enough to capture the large-size module and is enough to understand the way of traffic aggregation in the flow hierarchy. Another reason is that although the Infomap method can detect some “periphery nodes” (which in turn form a small-size module), such small-size modules are detected at each CL. Thus, the relation between the large-size module at the CL and the large-size module at lower-level CL is not suffered from the resolution limit.

5. Long-Term Change in Traffic Aggregation

Section 4 showed that the structure within a module can be represented as a star-like graph. It was also revealed that the structure in higher CLs has expanded by increasing the number of ASes within a module, whereas the structure in lower CLs has expanded with an increase in the number of modules. In this section, we use this structural analysis to investigate where the traffic will become concentrated. In particular, we focus on the traffic amount over intermodule links where large amounts of traffic are exchanged.

5.1. Relationship between Intermodule Links and Traffic Aggregation. The traffic concentration on links between modules is dependent on the structure of the Internet topology. In particular, the number of submodules influences the amount of traffic aggregated on the links between modules. This is because traffic aggregated inside each submodule is aggregated at an AS in a higher-level module, and the traffic aggregated at this AS is relayed via links between modules. We therefore investigate the number of submodules contained in a module. Figure 8 shows the average number of submodules contained in a module in each CL. The average number of submodules contained in a CL1 module increased until 2007, after which it can be seen to have slightly decreased. In levels below CL2, the average number of submodules has remained almost constant. In CL2, the average number of submodules

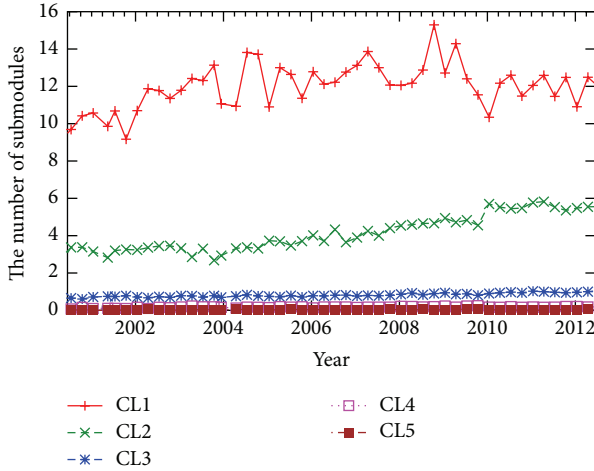


FIGURE 8: Average number of submodules contained in a module in each CL.

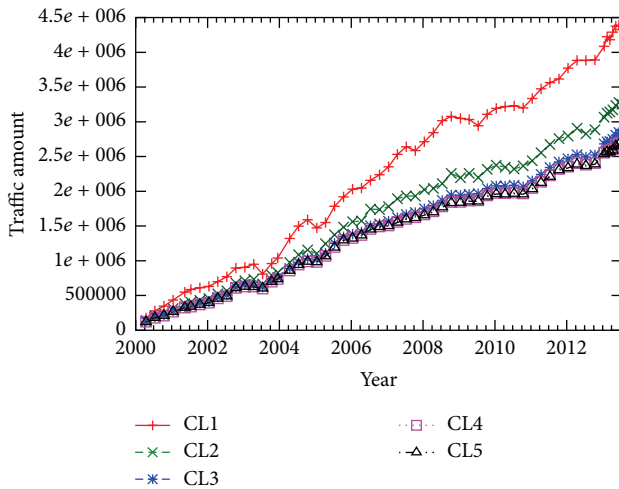


FIGURE 9: Average traffic amount on intermodule links in each CL.

has increased. The reason for this increase is that the number of CL3 modules has increased more than the number of CL2 modules, as shown in Section 4.4. Thus, more traffic has become concentrated on the links between CL2 modules.

5.2. Amount of Traffic Traversing Links between Modules. We now investigate the traffic concentration on links between modules. Figure 9 shows the increase in the average amount of traffic traversing links between modules in each CL. In obtaining the figure, we use (2) as a traffic demand between ASes and then calculate traffic amount of links. The average amount of traffic on links between CL1 modules has increased more than in other CLs. If this trend continues, more traffic will become concentrated on links between CL1 modules. The amount of traffic traversing links between CL2 modules also increased compared to the other CLs. In particular, the amount of traffic traversing links between CL2 modules in Figure 9 has slightly accelerated since 2011. The reason for the shift in 2011 may relate to the change of structure in

CL2 modules. In Figures 2(c) and 2(e), we can see that the increase in the assortativity and diameter of CL2 modules stopped around 2011. This implies that ASes having few links have tended to connect to an AS with the highest degree in a CL2 module after 2011. This trend leads to the increase in traffic aggregated on links between CL2 modules, and the acceleration in the amount of traffic on links between CL2 modules prevents the Internet from accommodating the overall increase in traffic. To determine how sensitive to γ our result is, we also examine by changing the value of γ from 238 at year 2004 to 3804 at year 2012 in Figure 15. From Figure 15, a similar tendency of traffic concentration is observed. By the traffic concentration, the operating and investment costs of routers increase. For example, the increase in processing cost leads to heatings problem and the power cost to cool routers, which is the primary contributor to an energy footprint, exponentially increases [45]. Moreover, an expansion of network equipment is needed according to the increase in the traffic volume. However, the transit fee that an AS receives from the other ASes does not increase more largely than the increase of traffic traversing the AS [5]. The traffic concentration will prevent ASes from continual maintenance and expansion of network equipment. Therefore, a new evolution process is needed to slow down the traffic concentration on links between CL1 modules and links between CL2 modules.

6. Evolution to Accommodate the Increase in Traffic Amount

Our analysis of the flow hierarchy shows that traffic is concentrated on links between CL1 modules and links between CL2 modules. Therefore, an evolution process that considers the global structure of the Internet topology is needed to slow down this increase in concentration. In this section, we examine a new evolution process that attempts to increase the number of links between lower-level modules to reduce the traffic concentration among higher-level modules. We explain our evolution process in Section 6.1 and then evaluate its performance in Section 6.2.

6.1. Evolution Process to Slow Down Traffic Concentration.

The results presented in Section 5 show that traffic has become increasingly concentrated on links between CL1 modules and links between CL2 modules. This is mainly because the number of ASes within CL1 and CL2 modules has increased, leading to an increase in the traffic generated in these modules. To continually accommodate the increase in traffic amount, the Internet topology requires a new evolution process to reduce this concentration at the links. Because the degree of traffic concentration on the links depends heavily on the global structure of the topology, our focus here is a global structure that can accommodate more traffic without increasing the concentration. For this purpose, we apply our evolution process in a centralized manner, rather than in the autonomous manner currently employed by ASes.

The basic approach of our evolution process is to construct more links between modules at lower CLs. With the links between lower CL modules, the traffic concentration

in the current Internet can be relaxed, as some of the traffic will no longer have to traverse links between higher-level modules. On the one hand, our evolution process is necessary to avoid traffic concentration among higher-level modules associated with the increase in traffic amount. On the other hand, our evolution process relies to some extent on the current topological characteristic that attempts to aggregate many paths into one link. In fact, the Internet topology has evolved such that a hub AS attracts more intramodule links (see discussion of Figure 2(b)). The hub AS aggregates and exchanges traffic from/to other modules. In the proposed evolution process, we must avoid traffic concentration among higher-level modules while retaining the characteristic of traffic aggregation used in the past. We therefore introduce a parameter w to represent the threshold of the number of links between hub ASes in different modules. As we increase w , the number of links between modules increases, which will lead to a relaxation in traffic concentration at higher-level modules. By changing the value of w , we are able to examine how the number of links between modules slows down the traffic concentration in links between higher-level modules. Formally, w is defined as follows. Let $E_v(\text{CL}x)$ be the set of links between a hub AS and another AS in a $\text{CL}x$ module, and let $E_h(\text{CL}x)$ be the set of links between $\text{CL}x$ modules. We define the ratio $R(\text{CL}x)$ of links in $E_h(\text{CL}x)$ to both $E_v(\text{CL}x)$ and $E_h(\text{CL}x)$ at $\text{CL}x$ as

$$R(\text{CL}x) = \frac{|E_h(\text{CL}x)|}{|E_v(\text{CL}x)| + |E_h(\text{CL}x)|}. \quad (4)$$

Then, our evolution process increases the links in $E_h(\text{CL}x)$ until $R(\text{CL}x)$ exceeds the threshold w . Figure 10 illustrates how $R(\text{CL}x)$ is calculated. A red node denotes a hub AS, which we call the gateway AS hereafter, in a module. A link between two red nodes is a link in $E_h(\text{CL}x)$, which is shown as a blue line. A link between a red node and a white node is a link in $E_v(\text{CL}x)$, which is shown as a red line. By increasing w , links in $E_h(\text{CL}x)$ are constructed between blue nodes. We then evolve a topology T using the following evolution process.

Step 1. Add new ASes to T .

Step 2. Add only one link for each new AS in T such that the new AS connects links to T .

Step 3. Calculate the flow hierarchy of T .

Step 4. Repeat the following steps from $\text{CL}6$ to $\text{CL}1$.

Step 4.1. Add a link between modules at the same CL .

Step 4.2. Calculate $R(\text{CL}x)$.

Step 4.3. If $R(\text{CL}x) < w$ and the connection among $\text{CL}x$ modules is not a full mesh, return to Step 4.1.

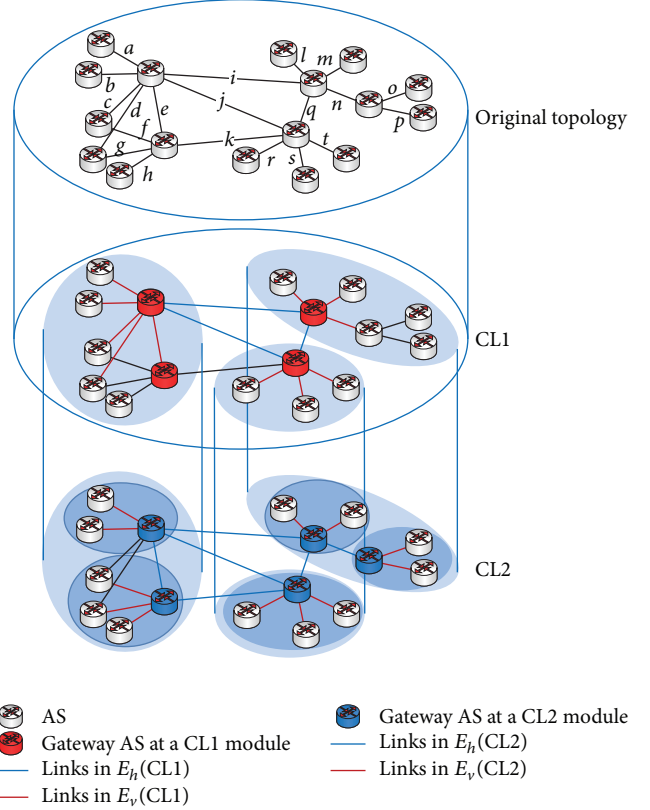


FIGURE 10: Illustrative example for $E_v(\text{CL}x)$ and $E_h(\text{CL}x)$; $E_h(\text{CL}1)$ is $\{a, b, c, d, e, l, m, n, r, s, t\}$ and $E_v(\text{CL}1)$ is $\{i, j, q\}$. $E_h(\text{CL}2)$ is $\{a, b, f, g, h, l, m, o, p, r, s, t\}$ and $E_v(\text{CL}2)$ is $\{e, i, j, k, n, q\}$. Thus, $R(\text{CL}1)$ is $3/14$ and $R(\text{CL}2)$ is $6/18$.

In Step 4.1, the link is constructed between gateway ASes, because a certain degree of traffic aggregation should be retained to preserve its characteristics.

Note that the current Internet does not have a mechanism which lets an AS know the location and the CL of the other gateway ASes. However, each AS can estimate whether an AS is a gateway from the AS paths in BGP tables. When most AS paths traverse a specific AS, the AS is considered as a gateway AS. Since the amount of traffic on links between modules in the same CL differs according to CL , as shown in Figure 9, the traffic amount on links connecting to gateway ASes also varies with the CL . By investigating the number of AS paths that traverse the gateway AS, the CL of the gateway AS can be estimated.

6.2. Effect of the Evolution Process

6.2.1. Backtracking the Internet Topology. We examine the effect of our evolution process in terms of slowing down the traffic concentration. For this purpose, we apply our evolution process to the Internet topology in the year 2000 and evolve the topology until 2013. Then, we compare the degree of traffic concentration in the evolved topology with that of the actual Internet topology in 2013.

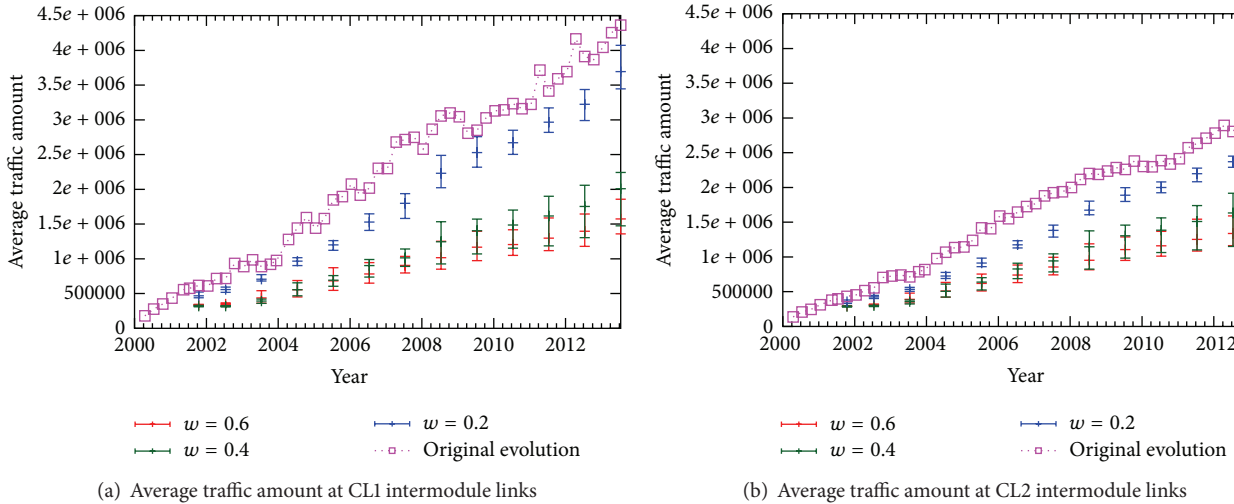


FIGURE 11: Average traffic amount at CL1 and CL2 intermodule links in the topology grown by applying the proposed evolution process.

To apply our evolution process, we first check the ASes and links added from year y to year $y + 1$ from $G(y)$ and $G(y + 1)$. Here, $G(y)$ represents the actual Internet topology at year y . We then evolve the topology such that the ASes are the same as those in the Internet topology in the next year. Links between ASes are constructed by the proposed evolution process described in Section 6.1. The evolution process is repeatedly applied 13 times; that is, the topology $T(y)$ is evolved to $T(y + 13)$. Note that when some ASes vanish at $G(y + 1)$, we remove them and their links from $T(y + 1)$ just after Step 1. If the topology $T(y + 1)$ becomes unconnected by this removal process, we select the largest connected component for further evolution. Selecting the largest connected component leads to a decrease in the number of ASes and links. However, we can confirm that the number of ASes in unselected connected components is less than 1% of all ASes, so the impact of this decrease is negligible.

At Step 3, we recalculate the flow hierarchy after adding links in Step 2 such that the flow hierarchy reflects the change of traffic aggregation altered by the link addition. At Step 4.1, we randomly select a pair of gateway ASes to construct a link on $E_h(CLx)$. Instead, we could calculate the optimal pair that minimizes the amount of traffic traversing links between higher-level modules. However, such a calculation is difficult in practice, because it requires complete information about the Internet topology and AS paths. Therefore, we randomly select a pair of gateway ASes and estimate the change in the amount of traffic traversing links. In this paper, we evolve the Internet 10 times with different random seeds and present the average change in the amount of traffic traversing links. After Step 4.1, when the number of links in $T(y + 1)$ is the same as the number of links in $G(y + 1)$, we stop applying our evolution process. After Step 4.3, we ensure that the number of links in $T(y + 1)$ is equal to that of $G(y + 1)$ for the purpose of comparison. We randomly select links from a set that is not included in $T(y + 1)$ but is included in $G(y + 1)$ and add the selected links to $T(y + 1)$. Finally, y is set to $y + 1$, and our evolution process is again applied until y becomes 2013.

6.2.2. Evaluation Results of the Proposed Evolution Process. To investigate how the number of links in $E_h(CLx)$ should be increased in the Internet topology, we evaluate the amount of traffic at links between higher-level modules in the topology evolved by our evolution policy. Figure 11 shows the average and the range of traffic amount on CL1 and CL2 intermodules links of 10 evolutions with different random seeds. The figures show results for $w = 0.2, 0.4$, and 0.6 . Note that the evolved topology has more links in $E_h(CLx)$ as we increase w . Figure 11 shows that the evolution policy slows down the increase in traffic at links between higher-level modules. When w is 0.2, this slowdown is small, because the size of $E_h(CLx)$ is small. In contrast, when the threshold w is set to 0.4 or 0.6, the slowdown effect is high. This is because traffic no longer needs to traverse links between higher-level modules. More importantly, the increase in traffic on links between CL2 modules has accelerated since 2011 in the original evolution, but this trend is not observed in Figure 11(b). We observe that the traffic concentration given by our evolution policy with $w = 0.6$ is not significantly different from that when $w = 0.4$. This suggests that when the size of $E_h(CLx)$ is above some threshold, the slowdown effect is not enhanced. We consider the traffic aggregated at links between higher-level modules to be adequately reduced when w is set to 0.4.

We finally investigate the influence of our evolution policy on the characteristics of the Internet. Figure 12 shows the average path length and the clustering coefficient of the Internet topology in January 2015 and the graphs evolved by our evolution policy. Average path length of topologies evolved by our evolution policy is larger than the Internet topology. However, when parameter w of our evolution policy is 0.4, the difference is marginal. Clustering coefficient of topology evolved by our evolution policy is lower than that of the Internet topology. However, since clustering coefficient of the Internet topology is also quite small, the difference of absolute value of clustering coefficient is small. These mean that

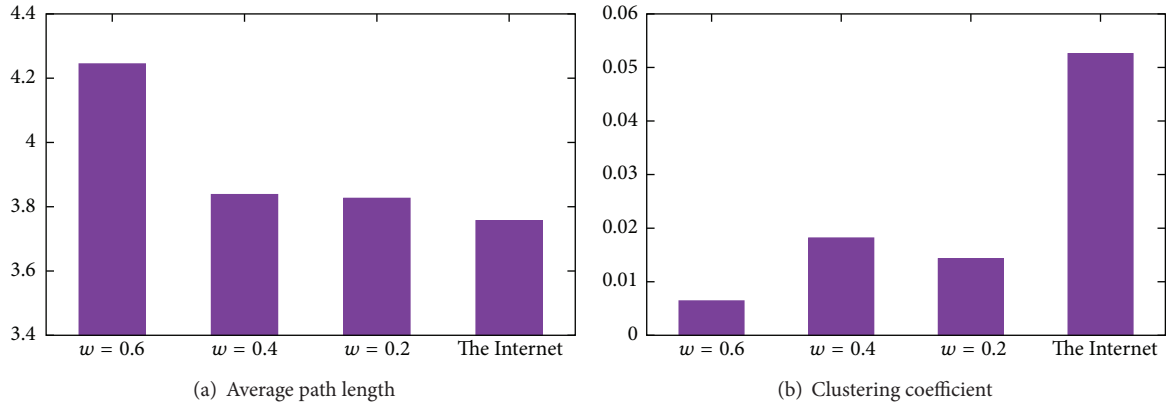


FIGURE 12: The influence on graph metrics by the evolution policy.

the influence of our evolution policy on the characteristics of the Internet topology is marginal.

These results mean that a suitable structure is derived when the threshold w is 0.4. Although our evolution policy with $w = 0.4$ slows down the traffic concentration, the volume of traffic on links between higher-level modules increases slightly. Therefore, there is a possibility that the traffic concentration will become a problem in the distant future. To further reduce the traffic concentration on these links, each AS exchanges information about which AS is a gateway in the modules at each CL. Thus, some feedback mechanism is required to achieve a suitable global structure and global performance. Under such a feedback mechanism, more suitable pairs of gateway ASes can be selected to construct links in $E_h(CLx)$. This evolution process may be difficult to realize in the current mechanism of link construction of ASes because this evolution process does not include the economic incentive for ASes. Our focus is not to develop a rigid evolution policy but investigate how the principles of evolution policy lead to the difference of the evolution of the global structure and whether it is possible to relax the future traffic concentration or not. Results show that our proposed evolution process can relax the traffic concentration on links between top-level modules by a half of the traffic concentration in the original evolution as shown in Figure 11. In practice, some economic incentives for promoting ASes to construct links based on the evolution process are necessary to optimize the performance of the global Internet, which is left for our future work.

7. Are Hyper Giants Necessary for the Evolution of the Internet?

Recently, the appearance of Hyper Giants, such as Google and Akamai, has impacted the traffic flow and evolution of the Internet topology. They generate huge amounts of traffic and send this across the Internet. Reference [14] found that the traffic amount sent by Hyper Giants is about 30% of the whole amount across the Internet, and the traffic amount generated and sent by Hyper Giants is expected to increase [5].

The appearance of Hyper Giants has influenced the structure of the Internet topology [5, 6, 8, 14, 46]. Hyper Giants construct peering links to ASes that use services provided by Hyper Giants, so that traffic sent by Hyper Giants does not traverse large ISPs. The primary reason that Hyper Giants construct a lot of peering links is to reduce the transit cost of traffic traversing large ISPs.

The increased number of peering links partly helps the Internet topology to achieve a suitable structure to continually accommodate an increase in traffic amount. This is because peering links are connected between modules, which are links in $E_h(CLx)$. However, the appearance of Hyper Giants alone will not allow the Internet topology to evolve sufficiently to accommodate the increase in traffic amount, because only traffic between a Hyper Giant and an ISP can be exchanged over the peering links. To accommodate the increase in traffic amount, some of the traffic aggregated at links suffering from overconcentration must traverse the other links. The peering links of Hyper Giants do not exchange traffic, but links between ISPs can. Therefore, it is important to consider not only the peering links of Hyper Giants, but also the connection among ISPs.

8. Conclusion

An evolution process that considers the global structure of the Internet topology is needed to accommodate future traffic amount. An analysis of the structure in the topology reveals where the traffic is concentrated, which enables us to develop an evolution policy to relax the overconcentration. Many works have shown that the Internet has a hierarchical structure [11–13]. Within this hierarchical structure, an AS aggregates traffic from lower-level ASes and relays the traffic to higher-level ASes. To identify the hierarchical nature of traffic aggregation, we investigated the long-term change in the structure of the Internet topology by analyzing the flow hierarchy. By examining the internal structure of a module, we found that each hub AS in a module is a gateway that aggregates and exchanges traffic from/to other modules. Furthermore, when the traffic demand is given by the gravity




FIGURE 13: The size of CLI modules that are divided by the Louvain method.

model, we showed that the amount of traffic traversing links between top-level modules and link between second-level modules has been rapidly increasing.

We considered a new evolution policy to avoid traffic concentration and then examined how this policy could slow down the traffic concentration compared with the actual evolution of the Internet topology. The basic approach behind our evolution policy is to construct more links between gateway ASes in different modules at the same level of the flow hierarchy, particularly at lower levels. While the topology retains the characteristic of traffic aggregation, a new policy is needed to avoid traffic concentration. To retain this characteristic, links between a gateway AS and other ASes in the same module should be preserved. We therefore introduced a threshold that determines the ratio of links between gateway ASes in different modules to the links between a gateway AS and other ASes. By varying this threshold, we examined how many links between gateway ASes are needed to slow down the traffic concentration. In evaluating the effect of our evolution policy, we found that the traffic concentration at links between higher-level modules decreased noticeably when the threshold was 0.4 or 0.6. We thus considered the traffic aggregated at links between higher-level modules to be adequately reduced when $w = 0.4$.

In future work, we will develop an evolution policy that considers the merits of each AS. Because the evolution of the Internet topology is not centrally controlled but an ensemble of individual link construction by each AS, the evolution policy should be applied to each AS. Indeed, [5, 47] investigated the evolution of the Internet topology from the viewpoint of game theoretic behavior by each AS. Future evolution policies must consider both the merit to individual ASes and the merit for the global structure of the Internet.

Appendix

A. The Impact of Resolution Limit on Analysis of the Evolution of the Flow Hierarchy

In Section 4, we analyzed the evolution of the Internet topology by investigating the evolution of the flow hierarchy. In the investigation, we used the Louvain method to exploit the flow hierarchy from the Internet topology. However, it is known that the Louvain method suffers from a resolution limit. The resolution limit is the characteristic scale of the smallest size of a module that the method can detect. To determine the effect of resolution limit on the analysis of the evolution of the flow hierarchy as shown in Table 5 and Figure 6, we analyze the evolution of each CL with the Infomap method [26], which does not much suffer from the resolution limit [44].

A.1. Analysis of the Size of Module by Infomap Method. We first investigated the size of module derived by the Louvain method and the Infomap method. Figure 13 shows the size of CLI modules derived by the Louvain method on 15 November 2013. x axis indicates the size of a module and y axis indicates the number of modules. The width of a bar is 2 in both of Figures 13(a) and 13(b). The number of modules containing fewer than 10 ASes is only about 10 and there are a few large modules containing more than 5000 ASes. Figure 14 shows the size of CLI modules derived by the Infomap method on 15 November 2013. x axis and y axis show the size of CLI modules derived by the Infomap method and the number of modules, respectively. By the Infomap method, much more amount of small-size modules appears compared to the Louvain method. This means that there is the effect of resolution limit on our analysis of the modular structure.

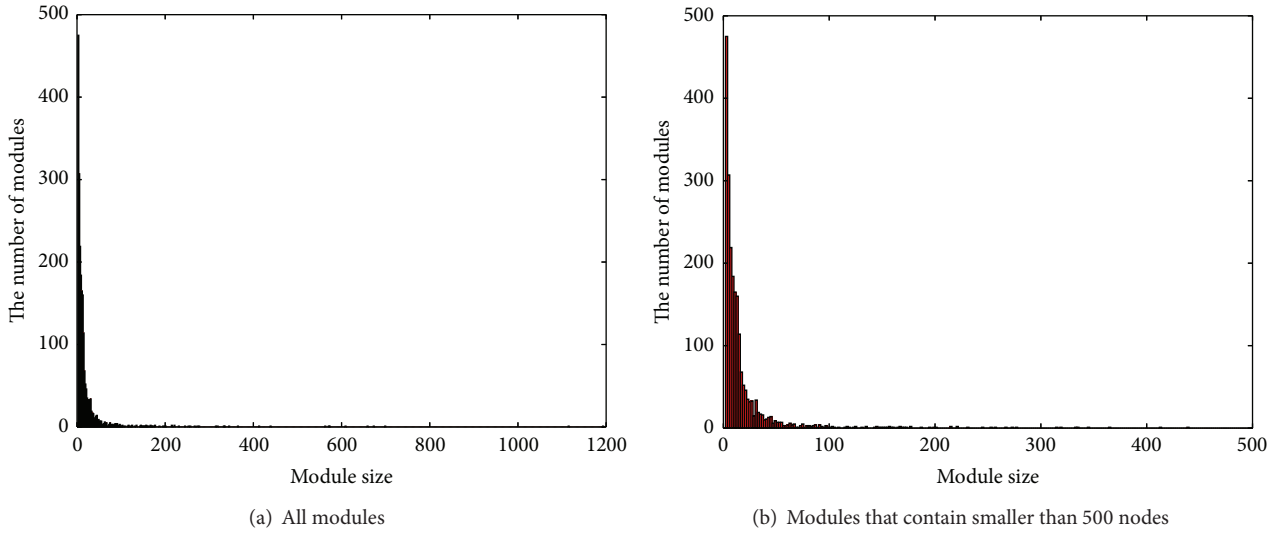
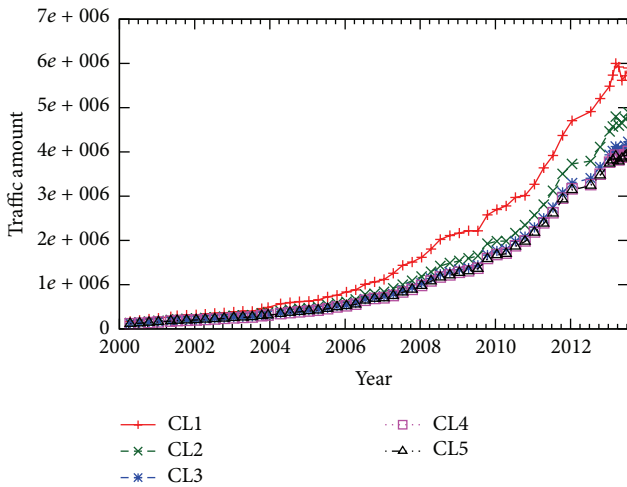


FIGURE 14: The size of CLI modules that are divided by the Infomap method.

FIGURE 15: The average traffic amount on intermodule links in each CL as the increase of γ .

A.2. Analysis of the Evolution of the Flow Hierarchy by Infomap Method. We next clarify whether our analysis of Table 5 and Figure 6 in our first submitted paper is affected by the resolution limit. Table 7 shows the number of modules in each CL derived by the Infomap method. The depth of the deepest module is always 3 from 2000 to 2012; that is, the depth has not changed. This result is the same as the result with the Louvain method. The number of modules in middle or bottom level such as CL2 and CL3 has more greatly increased than CL1 modules which is top level. This result also agrees with the result with the Louvain method.

Table 8 shows the average number of ASes in a module derived by the Infomap method. The increase of ASes in a CL1 module is slight. The reason is that there are a much more number of small-size CL1 modules as shown in Figure 14. Nevertheless, the number of ASes in CL1 modules

TABLE 7: The number of modules in each CL derived by the Infomap method.

Year	CL1	CL2	CL3
2000	445	573	4
2002	752	950	19
2004	887	1313	53
2006	1127	1625	25
2008	1379	2221	42
2010	1757	2604	49
2012	1991	3177	64

TABLE 8: The average number of ASes in a module derived by the Infomap method.

Year	CL1	CL2	CL3
2000	18.34	10.33	4
2002	18.1	9.16	5.95
2004	20.31	10.56	7.66
2006	20.34	10.87	4.52
2008	21.26	10.74	4.83
2010	20.26	10.91	17.67
2012	21.1	10.67	6.03

has increased compared to CL2 and CL3. This suggests that the number of ASes in large-size CL1 modules has greatly increased. From Tables 7 and 8, we consider that the structure in top level has expanded with the center pattern in Figure 6 since the structure in top level has expanded by increasing the number of ASes within a module. The structure in low level has expanded with the left pattern in Figure 6 since modules have increased more than top level. This result agrees with the result of analysis with the Louvain method.

Conflict of Interests

The authors declare that there is no conflict of interests regarding the publication of this paper.

References

- [1] University of Oregon, "Route Views project," 2013, <http://www.routeviews.org/>.
- [2] RIPE NCC, "RIS Raw Data," 2013, <http://www.ripe.net/data-tools/stats/ris/ris-raw-data>.
- [3] L. Atzori, A. Iera, and G. Morabito, "The internet of things: a survey," *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.
- [4] D. Miorandi, S. Sicari, F. De Pellegrini, and I. Chlamtac, "Internet of things: vision, applications and research challenges," *Ad Hoc Networks*, vol. 10, no. 7, pp. 1497–1516, 2012.
- [5] A. Dhamdhere and C. Dovrolis, "The Internet is flat: modeling the transition from a transit hierarchy to a peering mesh," in *Proceedings of the ACM 6th International Conference (Co-NEXT '10)*, pp. 1–12, December 2010.
- [6] A. Dhamdhere and C. Dovrolis, "Twelve years in the evolution of the internet ecosystem," *IEEE/ACM Transactions on Networking*, vol. 19, no. 5, pp. 1420–1433, 2011.
- [7] E. Gregoria, L. Lenzenib, and C. Orsini, "k-dense communities in the Internet AS-level topology," *Computer Networks*, vol. 57, pp. 1–10, 2013.
- [8] Y. Shavitt and U. Weinsberg, "Topological trends of Internet content providers," in *Proceedings of the 4th Annual Workshop on Simplifying Complex Networks for Practitioners (SIMPLEX '12)*, pp. 13–18, ACM, Lyon, France, April 2012.
- [9] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.
- [10] L. C. Freeman, "A set of measures of centrality based on betweenness," *Sociometry*, vol. 40, no. 1, pp. 35–41, 1977.
- [11] Z. G. Daniel, D. R. Figueiredo, S. Jaiswal, and L. Gao, "On the hierarchical structure of the logical Internet graph," in *Proceedings of the Scalability and Traffic Control in IP Networks*, vol. 4526, pp. 208–222, 2001.
- [12] R. Pastor-Satorras, A. Vázquez, and A. Vespignani, "Topology, hierarchy, and correlations in Internet graphs," *Complex Networks*, vol. 650, pp. 425–440, 2004.
- [13] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network topologies, power laws, and hierarchy," *Computer Communication Review*, vol. 32, p. 76, 2002.
- [14] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," in *Proceedings of the 7th International Conference on Autonomic Computing (SIGCOMM '10)*, pp. 75–86, September 2010.
- [15] M. E. J. Newman, "Modularity and community structure in networks," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 103, no. 23, pp. 8577–8582, 2006.
- [16] A. Elmokashfi, A. Kvalbein, and C. Dovrolis, "On the scalability of BGP: the roles of topology growth and update rate-limiting," in *Proceedings of the 4th ACM International Conference on Emerging Networking Experiments and Technologies (CoNEXT '08)*, pp. 1250–1261, ACM, Madrid, Spain, December 2008.
- [17] CAIDA, "The cooperative association for Internet data analysis," 2013, <http://www.caida.org/home/>.
- [18] S. C. Au, C. Leckie, A. Parhar, and G. Wong, "Efficient visualization of large routing topologies," *International Journal of Network Management*, vol. 14, no. 2, pp. 105–118, 2004.
- [19] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM '99)*, pp. 251–262, September 1999.
- [20] V. Rosato, L. Issacharoff, S. Meloni, D. Caligiore, and F. Tiriticco, "Is the topology of the internet network really fit to sustain its function?" *Physica A*, vol. 387, no. 7, pp. 1689–1704, 2008.
- [21] I. Poesse, B. Frank, B. Ager, G. Smaragdakis, S. Uhlig, and A. Feldmann, "Improving content delivery with PaDIS," *IEEE Internet Computing*, vol. 16, no. 3, pp. 46–52, 2012.
- [22] Y. Zhang, "Internet as-level topology archive," <http://irl.cs.ucla.edu/topology/>.
- [23] B. Zhang, R. Liu, D. Massey, and L. Zhang, "Collecting the Internet AS-level topology," *ACM SIGCOMM Computer Communication Review*, vol. 35, pp. 53–61, 2005.
- [24] B. Huffaker, M. Fomenkov, and K. Claffy, "Internet topology data comparison," CAIDA Technical Report, 2012.
- [25] R. Cohen, "The Internet dark matter: on the missing links in the AS connectivity map," in *Proceedings of the IEEE Conference on Computer Communications (INFOCOM '06)*, pp. 1–12, April 2006.
- [26] M. Rosvall and C. T. Bergstrom, "Maps of information flow reveal community structure in complex networks," *Proceedings of the National Academy of Sciences*, vol. 105, pp. 1118–1123, 2008.
- [27] A. Lancichinetti, F. Radicchi, J. J. Ramasco, and S. Fortunato, "Finding statistically significant communities in networks," *PLoS ONE*, vol. 6, no. 4, Article ID e18961, 2011.
- [28] V. D. Blondel, J. L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2008, no. 10, Article ID P10008, 2008.
- [29] M. Rosvall and C. T. Bergstrom, "Maps of random walks on complex networks reveal community structure," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 105, no. 4, pp. 1118–1123, 2008.
- [30] M. Roughan, A. Greenberg, C. Kalmanek, M. Rumsewicz, J. Yates, and Y. Zhang, "Experience in measuring backbone traffic variability: models, metrics, measurements and meaning," in *Proceedings of the 2nd ACM SIGCOMM Internet Measurement Workshop (IMW '02)*, pp. 91–92, November 2002.
- [31] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg, "Fast accurate computation of large-scale IP traffic matrices from link loads," in *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '03)*, pp. 206–217, June 2003.
- [32] H. Chang, S. Jamin, and W. Willinger, "To peer or not to peer: modeling the evolution of the Internet's AS-level topology," in *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM '06)*, pp. 1–12, April 2006.
- [33] H. Chang, S. Jamin, Z. M. Mao, and W. Willinger, "An empirical approach to modeling inter-AS traffic matrices," in *Proceedings of the 5th ACM SIGCOMM Conference on Internet Measurement (IMC '05)*, vol. 5, pp. 139–152, August 2005.
- [34] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger, "The origin of power laws in Internet topologies revisited," in *Proceedings of the IEEE INFOCOM*, pp. 608–617, June 2002.
- [35] D. Jiang and G. Hu, "GARCh model-based large-scale IP traffic matrix estimation," *IEEE Communications Letters*, vol. 13, no. 1, pp. 52–54, 2009.

- [36] G. Huston, "CIDR-report," 2013, <http://www.cidr-report.org/as2.0/>.
- [37] Cisco, "Cisco global cloud index: forecast and methodology, 2011–2016," 2012, <http://www.itworldcanada.com/assets/cisco-global-cloud-index-forecast-and-methodology-2011-2016>.
- [38] Cisco, "Cisco global cloud index forecasts cloud traffic to grow sixfold by 2016," 2012, <http://newsroom.cisco.com/release/1091855>.
- [39] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 733–745, 2001.
- [40] M. E. Newman, "Mixing patterns in networks," *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics*, vol. 67, no. 2, 2003.
- [41] R. Guimerà, M. Sales-Pardo, and L. A. N. Amaral, "Modularity from fluctuations in random graphs and complex networks," *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics*, vol. 70, no. 2, Article ID 025101, 2004.
- [42] S. Fortunato, "Community detection in graphs," *Physics Reports*, vol. 486, no. 3–5, pp. 75–174, 2010.
- [43] J. Reichardt and S. Bornholdt, "Statistical mechanics of community detection," *Physical Review E*, vol. 74, no. 1, Article ID 016110, 2006.
- [44] A. Lancichinetti and S. Fortunato, "Community detection algorithms: a comparative analysis," *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics*, vol. 80, no. 5, Article ID 056117, 2009.
- [45] X. Ma and K. Harfoush, "Traffic concentration for a green Internet," in *Proceedings of the 9th International Conference on High Capacity Optical Networks and Enabling Technologies (HONET '12)*, vol. 9, pp. 142–146, December 2012.
- [46] B. Krogfoss, M. Weldon, and L. Sofman, "Internet architecture evolution and the complex economies of content peering," *Bell Labs Technical Journal*, vol. 17, no. 1, pp. 163–184, 2012.
- [47] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "Peering strategy adoption by transit providers in the internet: a game theoretic approach," *ACM SIGMETRICS Performance Evaluation Review*, vol. 40, no. 2, pp. 38–41, 2012.

