Characterizing Inter-domain Rerouting by Betweenness Centrality after Disruptive Events

Yujing Liu, Xiapu Luo, Rocky K. C. Chang, and Jinshu Su

***Abstract*—Rerouting is not uncommon in nowadays Internet because it can be triggered by many root causes, such as network faults, routing attacks, etc. However, few methods effectively characterize rerouting in the whole Internet. In this paper, inspired by a well known network science metric - betweenness centrality, we propose a new approach to characterize inter-domain reroutings. By defining and analysing the variation of AS betweenness centrality for neighbouring-destination routes and global routes separately, our method empowers users to identify the temporal, topological, and relational characteristics of route changes. We apply our method to investigate the Internet's reactions to four different disruptive events, including Japan earthquake in March 2011, SEA-ME-WE 4 cable fault in April 2010, routing attack on YouTube in February 2008, and AS4761 hijacking event in January 2011. This examination reveals many new insights. For example, the route flapping and the congestion caused by the side-effect of rerouting after cable faults significantly degraded path quality. Moreover, direct providers of attackers and victims are the most critical positions for amplifying impact of prefix hijacking attacks. Such results shed light on how to implement effective reactions to network faults and how to deploy efficient defense mechanisms against routing attacks.**

***Index Terms*—Inter-domain routing, betweenness centrality, disruptive event, BGP.**

I. Introduction

NOWADAYS, the Internet has become a critical infrastructure of our society. It is of great significance to study the resilience of the Internet after disruptive events, such as natural disasters and manipulated attacks. From the perspective of network management, rerouting in the inter-domain routing system is a common quick-fix strategy to restore the affected communications. From the perspective of network security, the rerouting after an attack provides valuable first-hand information for examining the impact of the attack and designing resilient routing systems. Therefore, it is highly desirable to profile the reroutings and quantify their effects. However, few methods effectively characterize reroutings in the whole Internet. The lack of such knowledge hinders network administrators from taking further appropriate operations.

Previous research focused on analysing either the origins or the receivers of routing information. In fact, in the propagation process of routing messages, transmitters play a key role in shaping the impact of routing events. It is based on two observations. First, the unstable or malicious routing message will be widely spread by several crucial ASes (Autonomous Systems). Second, paths from different sources to different destinations which share common problematic transit ASes may show common behaviours. In network science, *betweenness centrality* is a very useful metric for measuring the importance of a vertex. It is motivated by the observation that an important vertex will show up in many shortest paths between two vertexes [1]. Likewise, in the inter-domain routing system, the more times an AS shows up in the AS paths, the more traffic it is responsible for transiting in the Internet. In this paper, we propose a transmitter-oriented metric *betweenness centrality of AS* to capture the receive-then-transmit nature of the Internet routing system. Since BGP (Border Gateway Protocol) is the de-facto standard inter-domain routing protocol, and its *AS path* attribute records AS-level routes of the Internet, we calculate AS betweenness centrality from publicly available BGP data in a particular period of time. Considering the different complexity of analysing neighbouring-destination routes and global routes, we develop different methods accordingly to exploit the full spectrum of this metric. These methods help identify several important characteristics of inter-domain rerouting, such as the time span when most routes changed, the ASes which were affected most, and the synchronization of route changes.

To demonstrate the usefulness of our methods, we apply them to real routing data related to four disruptive events. More precisely, we infer the cause of performance degradation on the monitored paths to Hong Kong after Japan earthquake in March 2011. The results show that the most impacted ASes along the paths were affected by a cable fault several hours after earthquake, and rerouting occurred frequently among the origin and the backup paths. Hence, we attribute the path-quality degradation to this unstable routing state. Next, we carry out a more fine-grained analysis at the IP- and geolocation-level to identify the root cause of the path performance degradation on the paths from Hong Kong to Europe in April 2010. It turns out to be the side effect of

Manuscript received 15 August 2012; revised 1 February 2013. This work is extended from "Characterizing Inter-domain Rerouting after Japan Earthquake" presented at IFIP Networking 2012. The work is supported by PCSIRT (Grant No. IRT1012); Program for Science and Technology Innovative Research Team in Higher Educational Institutions of Hunan Province (Network Technology, NUDT); Hunan Province Natural Science Foundation of China (11JJ7003); the National Natural Science Foundation of China (Grant Nos. 61070199, 61003303, 60903185, 61202396); the National High Technology Research and Development Program of China (Grant No. 2011AA01A103); a grant (ref. no. ITS/355/09) from the Innovation Technology Fund in Hong Kong; a grant (ref. no. H-ZL17) from the Joint Universities Computer Centre of Hong Kong; and a grant (ref no. G-YK26) from The Hong Kong Polytechnic University.

Y. Liu and J. Su are with School of Computer, National University of Defense Technology, Changsha, 410073 China (e-mail: {liuyujing,sjs}@nudt.edu.cn).

X. Luo and R. Chang are with Department of Computing, The Hong Kong Polytechnic University, Hong Kong S. A. R., China (e-mail: {csxluo,csrchang}@comp.polyu.edu.hk).

rerouting around the SEA-ME-WE 4's cable fault that led to substantial congestion to the backup paths. In the third case, we examine the impact of BGP prefix hijacking happened to YouTube in February 2008. The results enable us to identify the most efficient locations to deploy prevention and detection mechanisms against this attack, and evaluate the effect of recovery operations taken by ISPs. Finally, we analyse another hijacking event launched by AS4761 in January 2011. This event is different from the previous one since there are multiple topologically diversified victims involved.

Our method well captures the similarity and the difference of these four different types of network faults and routing attack cases. More precisely, the Japan earthquake entailed short-term instability whereas SEA-ME-WE 4 cable fault resulted in long-term congestion. During the YouTube routing attack, concentrated changes happened on the direct providers of the sole attacker and the victim, whereas in the AS4761 hijacking event, reroutings emerged on the common providers of multiple victims. Moreover, reroutings in the cases of network faults are topologically adjacent whereas those in the cases of prefix hijacking attacks are topologically diversified. These insights shed light on the classification of disrupt routing events, the design of effective reactions to network faults, and the efficient deployment strategy of defense mechanisms against routing attacks.

The paper is organized as follows. We present related works in Section 2 and introduce our methodology in Section 3. Then we apply it to characterize reroutings after Japan earthquake, SEA-ME-WE 4 cable fault, YouTube hijacking event and AS4761 hijacking event in Sections 4-7, respectively. Finally, we discuss their similarity and difference, and then conclude the paper in Section 8.

II. Related Work

Previous research on inter-domain routing analysed BGP dynamics from BGP updates. Li et al. classified dynamics into three categories: forwarding dynamics, policy fluctuations and pathological duplicates, which reflect different types of routing changes [2]. Xu et al. employed the principal component analysis (PCA) to transform and group BGP updates into different AS clusters affected by distinct major events [3]. Li et al. measured the impact of Internet earthquakes based on deviations from normal BGP update dynamics [4]. RIPE NCC, Renesys and BGPMon monitored BGP updates collected by their own infrastructures to study the impact on the Internet after network faults and BGP hijacking attacks [5], [6], [7].

Numerous, yet computationally intensive metrics, have been developed in these works. They can be classified into two categories: origin-oriented and receiver-oriented. The former one analyses routing information generated by each origin AS. However, it is inefficient for locating the root cause of network problems caused by the side effect of some disruptive events. For example, after 2011's earthquake, Cho et al. examined BGP messages generated by local ISPs in Japan and claimed that almost nothing unusual in BGP [8]. However, the earthquake did affect routing system of the Internet, such as our monitored paths from Europe, US, Australia and Japan to Hong Kong. The problem may occur on the way in between. Thus, the origin-oriented methods are insufficient and inefficient for identifying the root cause.

The latter category studies routing information received by monitors. Although such methods are useful for evaluating the scope of routing changes, their measuring ability is limited by the number and locations of vantage points. For example, although BGPMon reported the impact of the AS4761 hijacking event on a few ASes that received bogus prefix announcements [7], the real impact is not limited to the monitored ASes. The infectors will transmit the bogus messages to other ASes in the Internet and extend the infected scope. We believe that identifying the critical ASes which amplify the impact during the propagation process of routing information between origin and receiver is more important in studying possible defense mechanisms against hijacking attacks.

III. Methodology

Characterizing inter-domain rerouting is the first and foremost step to diagnose network faults and defend against attacks. The most essential characteristics are the answers to the following questions: 'When does the problem start and end? Where does it originate from? and How does it perform?'. Moreover, an efficient and effective metric is desired to quantify these characteristics. To answer these questions, we propose a transmitter-oriented metric - betweenness centrality of AS, which is a well known metric in network science but hasn't been utilized to explore routing changes before. Our analysis covers both neighbouring-destination routes and global routes, and employs statistics-based and PCA-based methods accordingly to characterize inter-domain rerouting from the temporal, topological and relational perspectives of the new metric. It is worth noting that this method can also be used to analyse other levels of networks other than AS if corresponding routing data is obtained.

A. Variance of AS betweenness centrality

We model the Internet as a graph $G = (V, E)$ where V is the set of all ASes, and E is the set of AS links. Let v be an AS in V, then its betweenness centrality $BC(v)$ is defined as

$$BC(v) = \frac{\sum_{u,w \in V} \sigma_{uw}(v)}{\sum_{u,w \in V} \sigma_{uw}} \ , \ u \neq w \neq v \qquad (1)$$

, where $\sigma_{uw}(v)$ denotes the total number of AS paths between u and w that pass through v, and σ_{uw} denotes the total number of AS paths between u and w. In Eqn. (1), we divide the value by the total amount of AS paths to normalize the betweenness centrality into [0, 1]. It's worth noting that in inter-domain routing system, AS paths are not necessarily the shortest paths among ASes. As a policy-based routing protocol, BGP allows each AS to choose its own routing polices in selecting the best route, announcing and accepting routes. The AS path must be policy-compliant, thereby follows certain patterns such as *valley-free* and *customer-prefer* [9].

Rerouting will result in the changes of some ASes' betweenness centrality. As shown in Eqn. (2), we define *variance of AS betweenness centrality* $\tilde{BC}_t(v)$ to quantify the difference

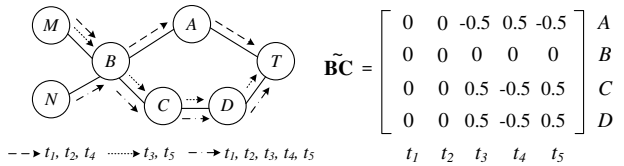

Fig. 1. An example of variance matrix.

of it measured at time $t-1$ and t. Positive value indicates extra AS paths are rerouted through this AS, while negative value indicates some paths are routed away from the AS. And the magnitude represents amount of paths.

$$\tilde{BC}_t(v) = BC_t(v) - BC_{t-1}(v) \quad (2)$$

We measure the variance of betweenness centrality of every AS in V for a period of time T and then obtain a *variance matrix* $\tilde{\mathbf{BC}}$, where every row represents a sequence of betweenness centrality changes associated with each AS (i.e., the element (i,j) in $\tilde{\mathbf{BC}}$ is equal to $\tilde{BC}_{t_j}(v_i)$ where $v_i \in V$ and $t_j \in T$).

We calculate $\tilde{\mathbf{BC}}$ by analysing BGP RIBs (Routing Information Bases) and updates collected by RIPE [10] and Route Views [11]. These two projects employ multiple Remote Route Collectors (RRCs) to establish BGP peering sessions with many ASes around the world, collect their routing information to all the other destination ASes, and periodically dump their BGP routing tables and updates. Therefore, we can get routes from those peering ASes to any other ASes in the Internet, and calculate variance of AS betweenness centrality from AS path of these routes. The available data from multiple vantage points reveals a broad and global though necessarily incomplete view of inter-domain routing over time. We use this data to sample the Internet's behaviour.

B. Characterize the rerouting for neighbouring-destination routes

Many of the network management and security tasks need to focus on the routes aiming at the same or neighbouring destinations. For these neighbouring-destination routes, we propose different statistical metrics to characterize rerouting from different dimensions of the variance matrix. Fig. 1 is a simple example illustrating routing changes from M and N to T on the left topology. The original routes are $M \to B \to A \to T$ and $N \to B \to C \to D \to T$. N doesn't change its route in the whole time period. At time slot t_3, M changes its route as $M \to B \to C \to D \to T$. Then it restores to the original one at t_4. At t_5, it changes again through C and D. The variance matrix of all transit ASes is shown on the right. Next, we will introduce the characterizing methods with this example.

1) Aggregated time of route changes: The variance matrix has two dimensions which are temporal and topological. In temporal dimension, every column of $\tilde{\mathbf{BC}}$ is a vector that contains every AS's betweenness centrality change at a certain time slot t, denoted as $\tilde{BC}_t =< \tilde{BC}_t(v_1), \tilde{BC}_t(v_2), ..., \tilde{BC}_t(v_n) >$, where $v_i \in V$. A larger variation indicates more route changes. We calculate the mean value of all ASes' absolute variation at time t, denoted as μ_t. Comparing all μ_t during an observed time window will show the time span having most route changes, which reveals the temporal characteristic of inter-domain rerouting. The μ_t value of the variance matrix in Fig. 1 is 0, 0, 0.375, 0.375 and 0.375 in turn. Apparently, reroutings happen from t_3 to t_5.

2) Worst affected components: The topological dimension of variance matrix includes all transit ASes in the Internet. Every row of $\tilde{\mathbf{BC}}$ is a vector that consists of betweenness centrality changes over time associated with a certain AS, denoted as $\tilde{BC}(v) =< \tilde{BC}_{t_1}(v), \tilde{BC}_{t_2}(v), ..., \tilde{BC}_{t_m}(v) >$, where $t_i \in T$. Higher diversity of this sequence of data indicates the AS experiences more unstable routing changes. Therefore, we use the standard deviation of the vector, denoted as δ_v, to measure the stability of AS v. Ranking all ASes with their δ_v can offer a list of members affected by the events in terms of descending instability. In our example, δ_v of A, B, C and D are 0.4183, 0, 0.4183 and 0.4183, indicating that A, C and D are equally unstable.

Large changes of betweenness centrality before and after a continuous time of instability indicate a quite different load distribution over the ASes. In other words, the routing system converges to a different state. Otherwise, the routing system just experiences unstable changing phase then goes back to the origin state. We define $\tilde{BC}_{(t_p,t_{p+q})}(v)$ as the changes of v's betweenness centrality between time t_p and t_{p+q}. It can be computed as $BC_{t_{p+q}}(v) - BC_{t_p}(v) = BC_{t_{p+q}}(v) - BC_{t_{p+q-1}}(v) + BC_{t_{p+q-1}}(v) - BC_{t_{p+q-2}}(v) + ... + BC_{t_{p+1}}(v) - BC_{t_p}(v) = \sum_{i=p+1}^{p+q} \tilde{BC}_{t_i}(v)$. For example, in Fig. 1's matrix, the $\tilde{BC}_{(t_3,t_5)}(v)$ of A, B, C and D are -0.5, 0, 0.5 and 0.5, which means that after continuous changes from t_3 to t_5, A loses paths passing through it, while C and D gain the same amount. These two aspects are considered to be the topological features of rerouting.

3) Synchronization of route changes: Investigating the relationship between routing behaviours from both temporal and topological aspects enables us to examine the synchronization of ASes' betweenness centrality changes, further more, to identify *correlative* and *backup* paths of original routes. Synchronization of routing changes indicates common causes beneath them. It is of great significance for the following study to separate these ASes.

We develop a divisive hierarchical clustering algorithm to analyse this synchronization. The basic idea is to evaluate the similarity between the patterns of any two ASes' changes. As shown in Fig. 2, the original data consists of all the vectors $\tilde{BC}(v)$. First, we split it into two clusters using K-Means algorithm with *absolute city block distance*, which is the sum of differences for all elements' absolute values in vectors (i.e., $\sum_{i=1}^{m} ||\tilde{BC}_{t_i}(v)| - |\tilde{BC}_{t_i}(v')||$). In each cluster, the betweenness centrality of every AS changes in similar magnitude at similar time. We assume that in an inter-domain routing event, only a few ASes act synchronously. So the smaller cluster contains the tightly related ASes. To validate this assumption, we examine the topology of these ASes, looking for clues to their relations. In order to further differentiate between correlation and backup, we partition

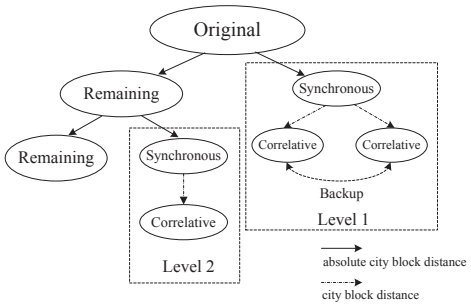

Fig. 2. Clustering algorithm for analysing synchronization of route changes.

these related ASes into two clusters again with just *city block distance* (i.e., $\sum_{i=1}^{m} |\tilde{BC}_{t_i}(v) - \tilde{BC}_{t_i}(v')|$). The intra-relation is correlative because the betweenness centrality of ASes within a cluster increases/decreases in similar quantity at same time, whereas inter-relation is backup because betweenness centrality increases in one cluster while decreases in the other at same time. Before differentiating them, we need to check whether or not there exist backup paths by calculating the intra-distance of correlative clusters and compare it with the intra-distance of their parent synchronous cluster [4]. If the difference is marginal and less than a threshold, it indicates that there is no apparent backup relation in the synchronous ASes. Therefore, we consider them all as correlative. The first isolated ASes and their routing behaviours are the dominant changes of the Internet's routing structure. We can keep clustering the remaining data to find different levels of changes as needed. For the simple example, A, C and D are sorted into one cluster because of their synchronous magnitude of variances. Furthermore, considering their changing directions in each time slot, we can infer that C and D are correlative whereas A is in a backup path to them.

C. Characterize the rerouting for global routes

In order to better understand inter-domain routing state of the entire Internet, we extend the analysis destinations to all routed IP prefixes in BGP. Unfortunately, our previous work shows that the above method for neighbour-destination routes is not suitable for examining global routes, especially for analysing synchronization of route changes [12]. It's because transit ASes employ different routing polices for different destinations. They don't show clear synchronous changing patterns globally. Moreover, the volume of original routing data is huge, while the correlations of ASes are complex. PCA (Principal Components Analysis) is a way of identifying patterns in data, and expressing the data to highlight their similarities and differences. We believe it's a suitable method to analyse global inter-domain reroutings.

PCA can transform the original space containing the observable variables into a new space of principal components, denoted as $\{PC_i\}, i = 1, ..., p$, which contain the variance inherent in the original data in descending order. The number of principal components p is the same as the number of original variables. Every principal component is a linear combination of the variables, and all of them are orthogonal. To capture $\theta\%$ of the variance of the original dataset, we find the smallest m such that $\frac{\sum_{i=1}^{m} \lambda_i}{\sum_{j=1}^{p} \lambda_j} \geq \theta\%$, where λ_i are the rank-ordered eigenvalues of the original covariance matrix. For more details about PCA, please refer to [13]. Here, we just focus on the relevant content of our study.

1) Dominant rerouting patterns: After applying PCA algorithm, we can get m principal components accounting for the most $\theta\%$ of the variances in the original variance matrix. These principal components can be considered as the dominant rerouting patterns among all the ASes. Concentrating on major changes in the global routing system is a better choice for analysing complicated and huge volume of routing data.

2) Head ASes in each pattern: As mentioned previously, every principal component is a linear combination of the variables. In the context of this paper, $PC_i = \sum_{v \in V} \alpha_{v,i} \tilde{BC}(v)$, where $\alpha_{v,i}$ is the coefficient (or *PC loading*) of $\tilde{BC}(v)$ for PC_i. It describes the contribution of $\tilde{BC}(v)$ to the variance captured by the ith principal component. So we sort loading values associated with all ASes in each dominant pattern and select the top 1 AS as the *head AS* in the corresponding rerouting pattern. The head AS is usually the worst affected AS in each pattern.

3) Propagation of rerouting in each pattern: The head AS can also be seen as the origin of routing changes. It is usually affected by the primary effect of the disruptive event. This change will propagate in the Internet within a scope of ASes which are affected by the secondary effect. We assume the related ASes will have similar rerouting patterns. Therefore, we propose to study this propagation by calculating the *Tanimoto similarity* of $\tilde{BC}(v)$ between every AS and the head AS in each dominant pattern. Tanimoto similarity of vectors X and Y is defined as

$$T(X,Y) = \frac{X \cdot Y}{|X|^2 + |Y|^2 - X \cdot Y} \tag{3}$$

It is a measure of the similarity between two vectors from both length and angle. The more similar the vectors are, the higher score they have. Sorting the similarity scores in descending order enables us to infer the propagation sequence of rerouting within a certain pattern. If the score is positive, it means the betweenness centrality of the two ASes increases/decreases in similar quantity at same time. We consider them as in the correlative paths. Meanwhile, if the score is negative, it means the betweenness centrality of one AS increases while the other decreases at same time, which implies they are in backup paths.

In addition, the topological locations of these ASes are also taken into account. Since routing changes only propagate in adjacent ASes rather than distant ones. We translate the topology graph G into a adjacent matrix **AD**. Element (i, j) in **AD** is equal to 1 if there is a link between v_i and v_j in E where $v_i, v_j \in V$. Otherwise, it's equal to 0. To see whether two ASes are adjacent, we can simply check whether the corresponding element in **AD** is 1. For the backup relations, two ASes are usually both adjacent to a third AS. In this case, we multiply **AD** with itself and check the corresponding element in the result matrix to see whether it's greater than 0. Because element (i, j) of **AD** $\times$ **AD** is calculated as

$\sum_{x=1}^{n} \mathbf{AD}(i,x) \times \mathbf{AD}(x,j)$, where $n = |V|$. It is actually the total number of adjacent ASes that v_i and v_j share. This topology examination is a crucial evidence to differentiate network faults and attacks, which will be explained later with study cases.

IV. Rerouting after Japan Earthquake

At 5:46 UTC on Mar. 11th, 2011, a massive 8.9-magnitude earthquake hit northeast Japan. Submarine cables connecting Japan to the rest of the world was damaged during the earthquake. As a result, the Internet relying on this underlying infrastructure was affected by this disaster. Our end-to-end Internet path measurement system OneProbe[1] [14] witnessed that 10 out of 19 monitored paths from Europe, US, Australia and Japan to HARNET (The Hong Kong Academic and Research NETwork) in Hong Kong suffered from high packet losses and increased RTTs during Mar. 11th and 12th. To diagnose the cause of this performance degradation, we seek for answers to these questions: Did any of the damaged cables affect those paths? Did the Internet reroute to bypass the fault segment? And how did the rerouting perform?

We collected BGP data during Mar. 9th to 12th, then divided them into a sequence of 10-minute time slots so that the routing changes could be distinguished in fineness of 10 minutes. By calculating every AS's betweenness centrality in each time slot, we constructed a variance matrix. We first limited the scope of our analysis to the routes from RRCs to HARNET's IP prefixes monitored by OneProbe. Results associated with these specific destinations helped us infer the causes of reported performance degradation. After that, we further analysed the routes for the entire routed IP prefix space in order to understand the global routing state.

A. Rerouting associated with HARNET

The variance matrix for routes associated with HARNET consists of 83 rows and 576 columns, which represent variances of betweenness centrality of 83 ASes in a period of 4 days. These ASes appear at least once in the AS paths of given routes.

Fig. 3 shows μ_t along with time. We use the first 2 days as a reference period, and calculate its average μ_t as a threshold to define normal range of routing changes. Therefore we can find the dominant period that contains the most 'abnormal' time slots after earthquake, which represents a continuous instability of routing state. The dominant period is from 9:00 to 16:00 on Mar. 11th, 3 hours later than the first wave of earthquake. This result is validated by Renesys [5]. They used the number of prefix outages as a metric, and inferred that network in Hong Kong was affected by follow-on events several hours later.

Fig. 4 shows ranked ASes with their standard deviations of vector $\tilde{BC}(v)$, i.e., δ_v. Higher values indicate more unstable routing states of ASes. To identify the worst affected components, we employ K-Means clustering algorithm to distinguish them from less affected ASes based on their δ_v.

[1]www.oneprobe.org

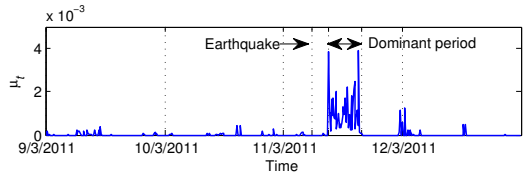

Fig. 3. Aggregated time of route changes with HARNET.

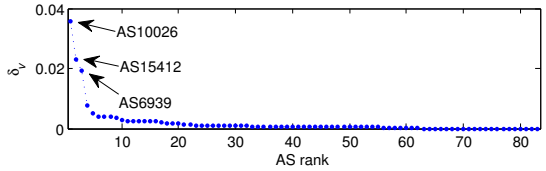

Fig. 4. Worst affected ASes with HARNET.

The first separated cluster contains 3 ASes which are AS10026 (PacNet), AS15412 (FLAG) and AS6939 (HURRICANE). All of them are major service providers of HARNET. This result implies the traffic passing through these 3 ASes shifts most frequently. According to TeleGeography's report [15], the Hong Kong-based cable-network operator PacNet reported damages to two segments of its East Asia Crossing submarine cable, which connects Japan to other parts of Asia. We believe this is the reason why AS10026's betweenness centrality varies noticeably after earthquake. Further analysis about relationships between ASes reveals the causes of changes in AS15412 and AS6939.

To investigate whether the cable fault-caused instability results in noticeable changes of load distribution, we examine $\tilde{BC}_{(t_b,t_e)}(v)$ of the 3 most unstable ASes, where t_b and t_e are the beginning and end time of the dominant period of rerouting. Then we compare them with the ASes' average betweenness centralities in the first 2 days. Their changing ratios are 0.32%, 0.68% and 0.75% separately, which indicates that although routes towards Hong Kong experience instability, few of them have changed after the unstable states.

We use the divisive hierarchical clustering algorithm to identify synchronous routing behaviours of ASes. Results show that level 1 cluster consists of the 3 worst affected ASes. AS10026 and AS6939 are correlative while AS15412 is a backup path of them. Given the inference that AS10026 is affected by the cable fault, we estimate that AS6939 is affected by the secondary effect of this cable damage via AS10026. In addition, AS15412 is a backup route for shifting traffic from AS10026 and AS6939. Level 2 consists of 4 ASes, in which AS22388 and AS7660 are backups of AS11537 and AS24167. Level 3 contains AS2914, AS4788 and AS6762, all of which are correlative. The hourly distributions of their variations in betweenness centrality are shown in Fig. 5. We stop at level 3 since changes are negligible after that. Fig. 6 shows the topology of AS paths from all RRCs to HARNET. The size of each AS is proportional to the log-value of its average betweenness centrality over the reference period (the first 2 days), which is considered as a metric of its importance in transmitting traffic for HARNET. It also serves as an evidence

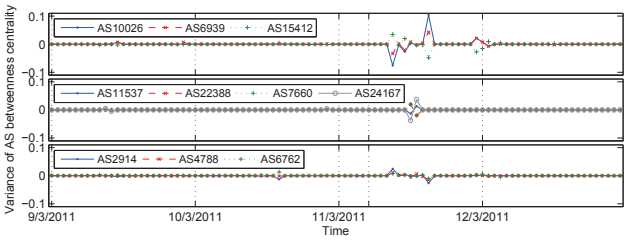

Fig. 5. Top 3 levels of synchronizations with HARNET.

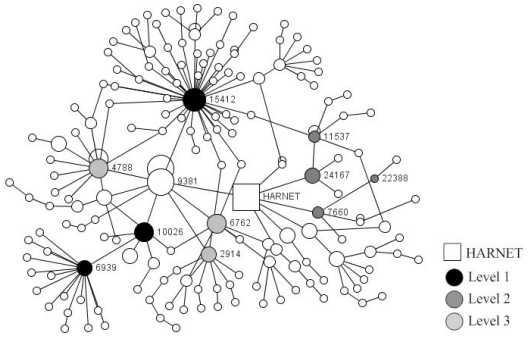

Fig. 6. Topology of AS paths from all RRCs to HARNET.

to infer HARNET's major providers. This figure validates our results in a topological view. For example, in level 1 cluster, AS10026 and AS6939 are neighbours while AS15412 and AS10026 are multi-hosts of HARNET's major provider AS9381. When one of the upstream network is damaged, switching to another provider is a quick solution.

Moreover, in Fig. 5, positive and negative variances of an AS alternate over time, which indicates flapping routes passing through the AS. We define flapping route as the route changes from one AS path to another, then changes back to the original one again, e.g. from AS path $A \to B \to C$ to $A \to D \to E \to C$ then back to $A \to B \to C$. Further statistic shows that there are 10534 route changes from Mar. 9th to 12th, 7821 of which are flapping routes, and 7071 happen in 11th. 103 of them flap more than 20 times. However, 96.9% of these routes are stick to the original one after flapping. A plausible explanation for route flapping is that the damaged cable segments entail decadent capacity of AS10026, then congestion of transmit traffic. Due to the co-location of data plane and control plane traffic of BGP, severe congestion can break down a BGP session between routers, forcing the traffic to be rerouted to backup paths. But after the congested traffic is shifted away, the previous failed session restores and the original route is utilized again. If this cyclical process repeats itself, there will be flapping routes in the inter-domain routing system [16]. However, the ground truth of this phenomenon should be explored further in future works.

B. Rerouting in the global Internet

The variance matrix for routes from RRCs to all prefixes consists of 5068 rows and 576 columns. These 5068 ASes are transit ASes which transmit traffic for other ones. The other more than 30,000 ASes are stub ASes which are not our targets for this study.

We set the threshold $\theta\%$ as 80%, and get 5 principal components as the dominant rerouting patterns in the global Internet. PC_1 accounts for as high as 62.3% of inherent variance of the original routing data. It indicates there exists an noticeable rerouting pattern after the earthquake.

After sorting the PC loadings of all $\tilde{BC}(v)$ for these 5 principal components, we identify the head AS in each dominant rerouting pattern. The results are listed in Table I. We are more interested in the first pattern, in which the head AS is still AS10026 (PacNet). It is evident that PacNet's regional cable faults have a vast impact on inter-domain rerouting of the Internet.

TABLE I
HEAD AS AND AFFECTED AS IN DOMINANT REROUTING PATTERNS

Dominant pattern	Head AS	Affected AS	Similarity
1	AS10026	AS6939	0.55
2	AS174	AS3257	-0.27
3	AS2914	null	0
4	AS38661	AS9457	0.94
5	AS3549	AS3257	0.24

Table I also shows the ASes which have a similar magnitude of more than 0.2 with the head AS in each dominant rerouting pattern. 0.2 is an optional threshold. We choose it with the thoughts of balancing the number of ASes and their similarity scores. The results indicate routing changes propagate among these ASes in different patterns. For instance, in the first pattern, AS6939 has a positive score of 0.55, and it is topological adjacent to AS10026. Hence we infer that AS6939 is in the correlative path with AS10026 in the global routing system, changing according to AS10026. It is affected by the side effect of PacNet's cable damage after the earthquake.

C. Possible reactive operations

According to our results, we infer that the performance degradation is mainly due to PacNet's cable faults 3 hours later after Japan earthquake. Paths passing through PacNet and HURRICANE change to FLAG as backup. The earthquake entails an unstable state of inter-domain routing system. A large number of inter-domain routes oscillate during that period of time. This instability of routing state is due to the current ad hoc rerouting mechanism. When the current route is no longer available, the router will choose a secondary-prefer route. The preference is often based on monetary and performance considerations. However, the present rerouting algorithm is not sufficient for maintaining the service quality in midst of disruptive events.

There are two possible reactive operations. The first one is to manually configure routing polices to help handle the network fault. This is a quick-fix solution after disruptive events. The key task of this operation is to locate and classify the root cause of reroutings along with path quality degradation. In this case, with the evidence explored from our method, we could cooperate with network operators in upstream ASes of PacNet to alternate routes, shifting some of Hong Kong-targeted traffic to other ASes instead of PacNet, in order to

bypass the damaged cable segments until its repair. Moreover, after adopting adjustment, it's important to evaluate effect of such operations on the Internet. The characteristics of rerouting measured by our method are also suitable for this task. The second reactive operation is to adjust the BGP algorithm. With input of real-time path performance and characteristics of routing behaviours, the route selection algorithm can be improved to maintain high service quality in the process of rerouting, performing an automatic feedback loop. The input data can be supported by our OneProbe measurement system and our characterizing method in this paper. We will explore it in future work.

V. Rerouting after SEA-ME-WE 4 Cable Fault

In April 2010, our measurement system OneProbe reported another cable fault. The South East Asia-Middle East-Western Europe 4 (SEA-ME-WE 4) submarine cable encountered a shunt fault on the Mediterranean segment on Apr. 14th, 2010 [17]. This event affected the forward path quality from Hong Kong to Europe. In addition to the AS-level characteristics of rerouting, we carry out a similar analysis in IP level and geolocation level. This fine-grained study helps us better understand how the routes were affected inside ASes. Moreover, geographical characteristics of rerouting enable us to analyse the physical topology of the Internet. From the output of this case study, we make an inference that a cable fault could significantly affect traffic on other non-faulty paths. The path performance could be improved if we are aware of this fact.

Along with measurement of data-path quality, we utilize Tcptraceroute for forward-path tracing. The data set for this study includes traces from Hong Kong to a BBC web server in the United Kingdom and a NOKIA web server in Finland from Apr. 13th to Apr. 16th, 2010. First of all, we map IPs to their ASes and do AS-level analysis. Fig. 7 (a) and (b) are RTT time series of the paths from Hong Kong to BBC and NOKIA, while (c) and (d) are aggregated time of route changes correspondingly. We can see that after cable fault on Apr. 14th, path quality declines on both paths, but routes don't show unusual changes. At 8:00 UTC on Apr. 16th, there is a peek slot of route changes on paths to BBC, after which the path performance is improved comparing with that of paths to NOKIA. Topological analysis shows that the significant route changes on the paths to BBC involve two ASes: AS15412 (FLAG) and AS6453 (TATA). The variance of their betweenness centrality at peek slot t_p, i.e. $\tilde{BC}_{t_p}(v)$, are -0.875 and 1, separately. Unlike the unstable reroutings after Japan earthquake, these changes are permanent. The above-mentioned evidences indicate that the cable fault affects paths through AS15412, leading to performance degradation. After changing from AS15412 to AS6453, the path quality is improved.

The AS-level analysis locates when and where the cable fault affects paths from Hong Kong to Europe. But it's too coarse to explain how. Hence we turn to IP-level analysis in these two critical ASes to gain more fine-grained insights. We put emphasis on the synchronization of IP path changes and apply our clustering algorithm. Fig. 8 shows top 3 clusters

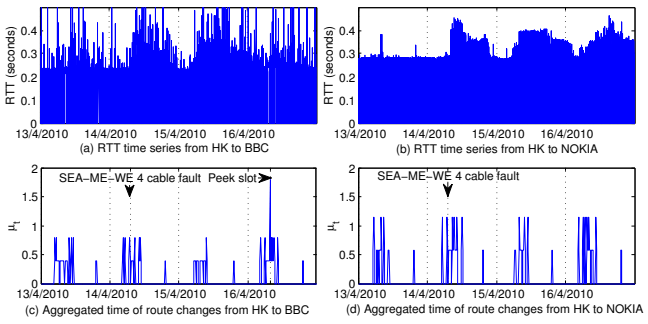

Fig. 7. Comparison of aggregated time of route changes after cable fault.

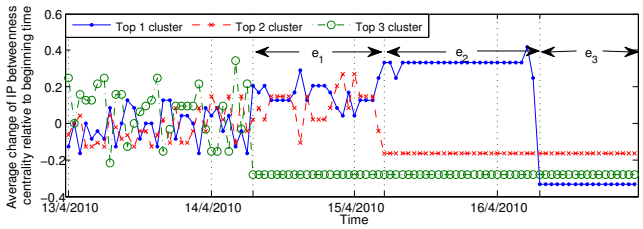

Fig. 8. Average changes of IP betweenness centrality of Top 3 clusters in AS15412.

of average changes of IP betweenness centrality relative to the beginning time, i.e. $\tilde{BC}_{(t_b,t)}(v)$, in AS15412. Within each cluster, the behaviours of IPs are correlative. Moreover, we resolve the IP hops' locations based on their DNS names, and find out that in cluster 1, IPs are located in Alexandria; in cluster 2, IPs are in Hong Kong and Alexandria; in cluster 3, IPs are in Mumbai and London. There are clearly three episodes of significant changes. In episode e_1, cluster 3 decreases while cluster 1 and 2 increase. In e_2, cluster 2 decreases whereas cluster 1 keeps increasing until episode e_3. In conjunction with previous geolocation analysis, we make a inference that subpaths in FLAG through Alexandria and Mumbai disappear episode by episode at the same time as path performance degradation. A possible explanation is that there are only two cable segments connecting landing points Alexandria and Palermo: FEA and SEA-ME-WE 4 [18]. After shunt fault on SEA-ME-WE 4, traffic was rerouted to FEA and congested this cable. Cable map of FLAG shows that they only use FEA cable for forwarding traffic between Hong Kong, Mumbai, Alexandria and London [19]. So the congestion on cable makes path performance decline.

We do the same IP-level analysis on AS6453 and find that after being rerouted to it, paths are passing through Hong Kong, Singapore, Chennai, Mumbai, Alexandria and London. TATA employs multiple cables to transit traffic. For example, between Alexandria and London, besides FEA, TATA also uses SEA-ME-WE 3 cable, providing diversity of underling infrastructures for routing system [20]. With this knowledge, we could improve path quality to NOKIA by collaborating with upstream AS of FLAG to change route towards other cable-rich or congestion-free ASes.

Our result of this case is consistent with previous research [17]. Comparing with the metric Jaccard distance, ours is a

vector of all the related components rather than a single value. So besides detecting problems, we are able to concentrate on the key ASes/IPs that need to be elaborate analysed so as to facilitate the locating process.

VI. Rerouting after YouTube being Hijacked

At 18:47 UTC on Feb. 24th, 2008, Pakistan Telecom (AS17557) advertised a route for prefix 208.65.153.0/24 to its provider PCCW Global (AS3491) by mistake. This is a sub-prefix of the network announced by YouTube (208.65.152.0/22). Due to longest prefix matching and forwarding, networks who receive this announcement will follow this more specific route to Pakistan Telecom instead of YouTube. The route propagated over the Internet and entailed a *sub-prefix hijacking* on a global scale. At 20:07 UTC, YouTube started reacting to this event, including announcing the same prefix as the hijacking one; announcing more specific routes (208.65.153.128/25 and 208.65.153.0/25); cooperating with PCCW Global to prepend Pakistan Telecom's AS number in the hijacking route and cooperating with PCCW Global to withdraw the hijacking route.

Prefix hijacking remains a major security threat to the inter-domain routing system. During the attack, an *attacker* announces the IP prefix which belongs to the *victim* network. Such bogus hijacking route propagates on the Internet and the ASes who accept the forged route become *infectors*. Data traffic from those infected ASes will be redirected to the attacker instead of the real destination, the victim.

Countermeasures against this attack include preventions before the attack, detections during the attack and reactions after the attack. Although many mechanisms have been proposed, few of them are practically deployed in a large scale. Inefficiency is the major issue. The actual routing behaviours after prefix hijacking are vital clues for us to propose efficient countermeasures. Taking the YouTube event as an example, we intend to answer the following questions: Instead of deploying in the whole Internet, where are the better locations to protect the routing system efficiently from being infected by Pakistan Telecom? What are the differences of the routing states before and after the prefix hijacking, which enable us to propose more efficient detections accordingly? And how effective are the reactions taken by YouTube after the event?

A. Rerouting associated with YouTube

Because this event only involved rerouting associated with a few prefixes, we apply the characterizing method for neighbour-destination routes. We collect BGP data during Feb. 24th, 2008, and then construct a variance matrix consisting of 66 rows and 144 columns.

By comparing μ_t and the timeline of the event, we find that before the hijacking, the routing state associated with YouTube is very stable. The aggregated time of route changes is right after the moment when Pakistan Telecom hijacked and YouTube reacted. Hence we divide the rerouting period into two phases - p_1 and p_2. p_1 lasted from 18:40 to 20:00, representing the hijacking phase. p_2 lasted from 20:00 to 21:20, representing the recovering phase. According to δ_v, the top 3 unstable ASes are AS3491 (PCCW Global), AS174 (Cogent) and AS3549 (Global Crossing). Moreover, their changes before and after p_1 and p_2 (i.e., $\tilde{BC}_{p_1}(v)$ and $\tilde{BC}_{p_2}(v)$) are also very prominent. After we apply our clustering algorithm, these 3 ASes show synchronous rerouting patterns, too. As shown in Fig. 9 (a), in p_1, AS3491 and AS174 are synchronous in opposite directions; likewise, in p_2, AS3549 and AS3491 are synchronous in opposite directions.

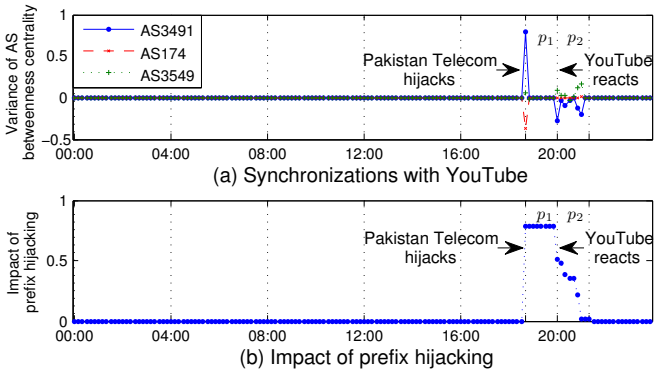

Fig. 9. Relationship between synchronizations and impact of attack with YouTube.

B. Insights of defense mechanisms against prefix hijacking

The unique routing changes of these 3 ASes in the hijacking and the recovering phases attract us to investigate them carefully. AS3491 is the direct provider of the attacker - Pakistan Telecom. Whereas AS174 and AS3549 are the direct providers of the victim - YouTube. Direct providers of an AS are the adjacent upstream ASes providing Internet services. We get the AS relationship information from BGP routing data and CAIDA's data source [21]. From the relationship of their reroutings and the impact of prefix hijacking as shown in Fig. 9, we get additional insights of defense mechanisms against this attack.

In an IP prefix hijacking, the more infectors there are, the more data traffic will be redirected and the greater impact an attack will have. In this paper, we define the impact of a prefix hijacking attack as the percentage of infected ASes in the whole Internet. Fig. 9 (b) shows the change of YouTube event's impact on Feb. 24th, 2008. We identify infectors by analysing MOAS (Multi Origin AS) conflict in BGP data during that day.

In p_1, right after Pakistan Telecom announces the hijacking route, the betweenness centrality of AS3491 increases up to 0.79. Then, the impact of this attack rises to the same value 0.79. It indicates that all the hijacking routes infecting ASes in the Internet contain AS3491. This attacker's direct provider is critical for propagating forged routes and amplifying the impact. If it has deployed some prevention mechanisms such as defensive filtering, it could filter out this bogus route which announces an IP prefix that doesn't belong to its customer. Inspired by this evidence, we infer that letting direct providers prevent their customers from launching hijacking is a more

efficient method than deploying preventions on every ASes to fight against any attackers in the Internet.

At the same time, the betweenness centrality of AS174 decreases dramatically. This is resulted from attacker's announcing more attractive routes to YouTube. Lots of original paths passing through AS174 to YouTube shift to other ASes. This sharply change on the direct providers of victim is a vital clue for the customer being hijacked. Moreover, the bigger impact an attack has, the greater change the direct providers experience. We believe it's also of great significance for deploying more efficient detection mechanisms. Detections employing this evidence only need to be deployed on transit ASes in the Internet, which are much fewer than stub ASes. In addition, the targets of detection are limited into the deployers' customers.

During p_2, YouTube takes a series of actions to restore routing state. After the same prefix as the hijacking one was announced, the impact of attack was reduced by 0.30. With two identical prefixes in the routing system, Pakistan Telecom continued attracting some of YouTube's traffic. After the longer prefixes were announced, the impact was further reduced by 0.13. After AS3491 prepended Pakistan Telecom's AS number in hijacking route, the impact was reduced by 0.13, too. It's because longer AS path makes the hijacking route less attractive. After AS3491 withdrawed the hijacking route, the impact was reduced to 0. This reaction finally stopped the hijacking of YouTube. During this process, the betweenness centrality of AS3491 changes exactly according to the changes of this event's impact, which is another evidence of its importance on prevention, detection and reaction of this attack. After restoration, the distribution of AS betweenness centrality changes. AS3549, another direct provider of YouTube attracted more routing paths to pass through it.

VII. Rerouting after AS4761 Hijacking Event

Another prefix hijacking event happened on Jan. 14th, 2011 [7]. AS4761 originated approximately 2800 new unique prefixes belonging to more than 800 ASes. These abnormal announcements were launched between 12:19 and 12:57 UTC. We take this event as another example because there are multiple topologically diversified victims. Analysis results of our method provide different yet important insights to this type of hijacking.

We parse BGP data from 12:00 to 14:00 that day. The aggregated time of route changes is from 12:10 to 12:50. The worst affected ASes are included in two clusters, indicating two major routing patterns. Fig. 10 shows the behaviour of 'head AS' on opposite directions in each cluster, and their relationship with impact of the event. AS9505 and AS7018 are in top 1 cluster, while AS3491 and AS7473 are in top 2 cluster. The hijacking period is from 12:10 to 12:20, in which the betweenness centrality of AS9505 increases sharply and that of AS7018 decreases. The recovering period is from 12:20 to 12:50 when impact of hijacking declines step by step. During this period, AS9505 decreases and AS7018 increases. AS3491 and AS7473 show similar routing pattern 10 minutes later than ASes in cluster 1, making a slow recovering process.

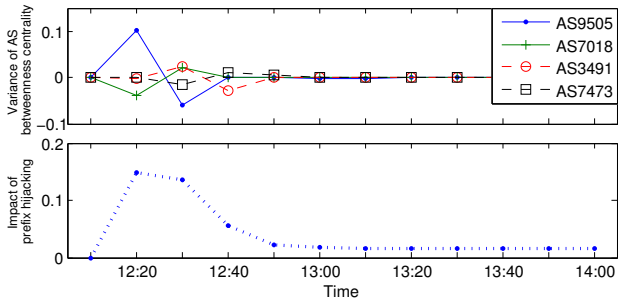

Fig. 10. Relationship between synchronizations and impact of AS4761 hijacking.

We elaborate more on how we could leverage the results to defend against hijacking with top 1 cluster. AS9505 is the direct provider of the 'attacker' AS4761 (it may launch the prefix hijacking due to misconfiguration, here we refer it as 'attacker' according to its behaviour, not its motivation). It's the key factor to amplify the impact of attack. So we make a similar conclusion as YouTube event that it's efficient to deploy prevention mechanisms on the direct providers to prohibit their customers from launching attacks.

For detection, the abnormal increase of betweenness centrality on AS9505 and the decrease of betweenness centrality on AS7018 shows an evidence of potential hijacking. In this case, a large amount of victims spread around the Internet. AS7018 is the common provider of most victims. In-depth analysis shows that 90.3 % victims are customers of AS7018, and 55.1 % are direct customers. This gives us a hint that deploying detections on common but less providers is a trade-off between accuracy of detections and cost of deployments. We could find an appropriate balance from deployed on all the transit ASes to a dozen of tier-1 ASes which are the top-tier providers of all ASes. Moreover, the synchronization of the two ASes in opposite directions is not due to the rerouting between backup paths. We draw this conclusion according to the analysis of their topological relationship with victim ASes. Only 2 victims are their common customers, providing a more affirmative indication that this event is a routing attack rather than a normal rerouting.

Similarly, for reaction, it's efficient for AS9505 to withdraw the hijacking routes and for AS7018 to re-announce more attractive routes towards victims. There is only one attacker in this event. It's rational to infer that if multiple attackers exist in a prefix hijacking, their common providers are practical positions for prevention, detection and reaction.

VIII. Discussion and Conclusion

In this paper, we propose a network science-inspired metric, the variance of AS betweenness centrality, to characterize behaviours of every transit AS in the route propagation process. Their temporal, topological and relational characteristics empower us to better understand the dynamics of inter-domain reroutings after four disruptive events. From the results, we identify the similarity and the difference between these two network fault cases and two BGP prefix hijacking attacks.

These insights are useful to improve the Internet fault and security management.

Reroutings after Japan earthquake are temporary and flapping changes. The performance degradation of monitored paths is attributed to the unstable routing state. While after SEA-ME-WE 4 cable fault, the permanent AS-level route changes improve path quality. Fine-grained IP-level analysis indicates an uncongested cable was affected by the side-effect of the reroutings around cable fault. With the knowledge of root cause and origin of rerouting, we could carefully adjust the routing configuration to bypass network fault while keeping the path performance.

The BGP prefix hijacking attack on YouTube has only one attacker and one victim. Our investigation shows that the direct providers of the attacker and the direct providers of the victim are in critical positions to deploy efficient prevention, detection, and reaction mechanisms. The AS4761 hijacking event has one attacker and multiple victims. In this case, the variance of AS betweenness centrality reflects emerging changes of common reroutings. So the common providers of victims become prominent. The differences between these two events provide hints to balance the scope of protection and the scale of deployment.

Furthermore, reroutings of network faults are usually topologically adjacent or geographically clustered, whereas reroutings after prefix hijackings usually synchronize on topologically diversified ASes. The difference of these two types of disruptive events shed lights on how to classify root causes and conduct proper reactive operations.

In this paper, we put emphasis on the betweenness centrality of vertexes in a network. However, the further study of links will certainly provide us more and deeper insights of the characteristics of reroutings in the Internet, especially in the scenarios where there is high likelihood of cable faults. In future work, we will examine other perspectives of the centrality metric and try to further understand and improve the inter-domain routing system by exploring the advances of network science.

REFERENCES

[1] L. da F. Costa, F. A. Rodrigues, G. Travieso, and P. R. Villas Boas, "Characterization of Complex Networks: A Survey of measurements," *Advances in Physics*, vol. 56, no. 1, pp. 167-242, 2007.

[2] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz, "BGP Routing Dynamics Revisited," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 2, pp. 7-16, 2007.

[3] K. Xu, J. Chandrashekar, and Z. L. Zhang, "Inferring Major Events from BGP Update Streams," Technical Report 04-043. http://www.cs.umn.edu/research/technical_reports.php?page=report&report_id=04-043, 2004.

[4] J. Li, and S. Brooks, "I-seismograph: Observing and Measuring Internet Earthquakes," In *Proc. IEEE INFOCOM*, 2011.

[5] Renesys blog, "Japan Quake," http://www.renesys.com/blog/2011/03/japan-quake.shtml

[6] RIPE news, "YouTube Hijacking: A RIPE NCC RIS case study," http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study.

[7] "'Hijack' by AS4761 - Indosat, a quick report," http://www.bgpmon.net/hijack-by-as4761-indosat-a-quick-report/.

[8] K. Cho, C. Pelsser, R. Bush, and Y. Won, "The Japan Earthquake: the Impact on Traffic and Routing Observed by a Local ISP," In *Proc. ACM SWID*, 2011.

[9] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 733-745, 2001.

[10] RIPE RIS, http://www.ripe.net/data-tools/stats/ris/routing-information-service.

[11] University of Oregon Route Views Project, http://www.routeviews.org/.

[12] Y. Liu, X. Luo, R. Chang, and J. Su, "Characterizing Inter-domain Rerouting after Japan Earthquake," In *Proc. IFIP Networking*, 2012.

[13] I. T. Jolliffe, *Principal Component Analysis (2nd edition)*, Springer Series in Statistics, 2002.

[14] X. Luo, E. Chan and R. Chang, "Design and Implementation of TCP Data Probes for Reliable and Metric-Rich Network Path Monitoring," In *Proc. USENIX Annual Tech. Conf.*, 2009.

[15] "Why the Japan Earthquake Didn't Take Down the Country's Internet," http://spectrum.ieee.org/tech-talk/telecom/internet/why-the-japan-earthquake-didnt-cripple-the-countrys-internet.

[16] M. Schuchard, E. Y. Vasserman, A. Mohaisen, D. F. Kune, N. Hopper, and Y. Kim, "Losing Control of the Internet: Using the Data Plane to Attack the Control Plane," In *Proc. CCS '10*, 2010.

[17] E. Chan, X. Luo, W. Fok, W. Li, and R. Chang, "Non-cooperative Diagnosis of Submarine Cable Faults," In *Proc. PAM*, 2011.

[18] TeleGeography - Submarine Cable Map, http://www.submarinecablemap.com.

[19] Reliance Globalcom, http://www.relianceglobalcom.com/RGCOM_CoverageMap.html.

[20] Global footprint map - Tata Communications, http://www.tatacommunications.com/map/gfp.html.

[21] CAIDA AS Relationships, http://www.caida.org/data/active/as-relationships/.

Yujing Liu received her M.Sc. degree in Computer Science from the National University of Defense Technology, China, in 2009. She is now a Ph.D. candidate in the School of Computer, National University of Defense Technology, China. She spent one year as a visiting student in the Hong Kong Polytechnic University in 2011. Her research interests include Internet routing and network security.

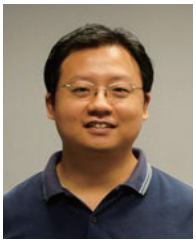

Xiapu Luo received his Ph.D. degree in Computer Science from the Hong Kong Polytechnic University in 2007 and then spent two years at the Georgia Institute of Technology as a post-doctoral research fellow. He is now a research assistant professor in the Department of Computing at the Hong Kong Polytechnic University. He is also a researcher affiliated with the Shenzhen Research Institute of the Hong Kong Polytechnic University. His current research focuses on network security and privacy, Internet measurement, and Smartphone security.

Rocky K. C. Chang received his Ph.D. degree in Computer Engineering from Rensselaer Polytechnic Institute in 1990. Immediately after that, he joined the IBM Thomas J. Watson Research Center working on performance analysis and simulation tools till July 1993. He then joined the Department of Computing at The Hong Kong Polytechnic University in 1993, where he is now an associate professor. His research interests include Network Measurement Systems, Internet Infrastructure Security and Privacy, and QoE Evaluation of Network Services.

Jinshu Su received his B.S. degree in Mathematics from Nankai University, Tianjin, China, in 1985 and the M.Sc. and Ph.D. degrees in Computer Science from the National University of Defense Technology, Changsha, China, in 1988 and 2000, respectively. He is a professor in the School of Computer, National University of Defense Technology. His research interests include Internet architecture, Internet routing, security, and wireless networks.