

Characterizing Inter-domain Rerouting after Japan Earthquake

Yujing Liu, Xiapu Luo, Rocky Chang, Jinshu Su

► **To cite this version:**

Yujing Liu, Xiapu Luo, Rocky Chang, Jinshu Su. Characterizing Inter-domain Rerouting after Japan Earthquake. Robert Bestak; Lukas Kencl; Li Erran Li; Joerg Widmer; Hao Yin. 11th International Networking Conference (NETWORKING), May 2012, Prague, Czech Republic. Springer, Lecture Notes in Computer Science, LNCS-7290 (Part II), pp.124-135, 2012, NETWORKING 2012. <10.1007/978-3-642-30054-7_10>. <hal-01531953>

HAL Id: hal-01531953

<https://hal.inria.fr/hal-01531953>

Submitted on 2 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Characterizing Inter-domain Rerouting after Japan Earthquake

Yujing Liu¹, Xiapu Luo², Rocky K. C. Chang², and Jinshu Su¹

¹ School of Computer, National University of Defense Technology
{liuyujing,sjs}@nudt.edu.cn

² Department of Computing, The Hong Kong Polytechnic University
{csxluo,csrchang}@comp.polyu.edu.hk

Abstract. The Internet is designed to survive when some of its components become faulty. Rerouting is an important approach to bypassing faulty links. In this paper, we propose a method to characterize inter-domain rerouting as a result of the massive earthquake in Japan on March 2011. Moreover, the characterization helps correlate inter-domain rerouting events and end-to-end path-quality degradation measured in Hong Kong. We analyse the variation of AS betweenness centrality extracted from BGP data to identify the time span when most routes changed, the ASes that were most seriously affected, and the correlative and backup paths. The results show that three major providers of inbound traffic to Hong Kong were affected by unstable routing state caused by a cable fault after the earthquake. Our work provides a new method of utilizing control plane's information to diagnose data plane's problem.

Keywords: inter-domain routing; betweenness centrality; Japan earthquake; BGP

1 Introduction

At 5:46 GMT on Mar. 11th, 2011, a massive 8.9-magnitude earthquake hit north-east Japan, causing dozens of deaths, more than 80 fires, and a 10-meter tsunami along parts of the country's coastline. Nuclear power plant and Japan's largest electric utility were taken down. Fiber optic network of undersea cables that connect Japan to the rest of the world was damaged when the earthquake struck beneath the Pacific seafloor [1]. As a result, the Internet relying on this underlying infrastructure to deliver data traffic was affected by this disaster. Many telecom operators reported service disruptions and performance degradations in the communications with either Japan or other places around the world.

Our end-to-end Internet path measurement system OneProbe¹ [2] 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

¹ www.oneprobe.org

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?

In this paper, we propose a new approach to characterize inter-domain rerouting in the AS level after network events. The Internet is highly distributed and autonomous in nature. Its routing state is profiled by the behaviour of every component in it. According to the evidence that rerouting entails changes of traffic load distribution, we measure the load variance of every AS in terms of AS betweenness centrality to depict properties of rerouting in the Internet. Since Border Gateway Protocol (BGP) 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 RIBs and updates collected by RIPE [3] and Route Views [4] in a particular period of time. By exploiting different perspectives of this metric, we can identify several important characteristics of inter-domain rerouting such as the time span when most routes changed, the ASes which were affected worst, and the correlative and backup paths. We apply our methodology to infer the cause of performance degradation on the monitored paths to Hong Kong after Japan earthquake. The results show that the worst impacted ASes along the paths are affected by a cable fault several hours after earthquake, and rerouting occurs frequently among origin and backup paths. Hence, we attribute the path-quality degradation to this unstable routing state. Moreover, we investigate how the Internet performed in a global view after the quake. The study of this event serves as a practical case of leveraging control plane’s information to diagnose data plane’s problem.

The paper is organized as follows. We introduce our methodology in Sec. 2. Then we apply it to characterize inter-domain rerouting after Japan earthquake in Sec. 3. After discussing the related works in Sec. 4, we conclude the paper and present future works in Sec. 5.

2 Methodology

Many researches study inter-domain routing behaviour by analysing BGP dynamics for different purposes, such as improving the Internet functioning, troubleshooting BGP, detecting anomalies and monitoring impact of disruptive events on the Internet [5–8]. Numerous yet computationally intensive metrics have been developed to achieve the goals. Most of them are prefix oriented (origins of routing information) whereas some explore the point of observation (receivers of routing information). In this paper, with the novel attempt to diagnose data plane’s problem, we focus on the changes in load carried by every AS in-between to investigate how different routing systems react after network events. To assess the load of an AS, we propose utilizing betweenness centrality as a metric that was introduced in social networks first [9, 10]. In this section, we will explain how to identify temporal, topological and relational characteristics of inter-domain rerouting from different perspectives of this single metric.

2.1 Variance of AS betweenness centrality

In networks, the more shortest paths a vertex or edge participates in, the more important this vertex or edge is. Betweenness centrality is a metric to quantify the importance of a vertex or edge in a network [11]. In this paper, we employ it to measure how many AS paths pass through a certain AS. Since the more AS paths an AS participates in, the more traffic the AS will transmit. Therefore, we consider it as a proper metric to evaluate the load on such AS. 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 the betweenness centrality $BC(v)$ is defined as

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

, where the summation is over all pairs of ASes such that $u \neq w \neq v$. $\sigma_{uw}(v)$ 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 . 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 policies in selecting the best route, announcing and accepting routes. The AS path from source to destination must be policy-compliant, thereby follows certain patterns such as *valley-free* and *customer-prefer* [12].

Inter-domain rerouting will result in 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 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.

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

If we measure the variance of betweenness centrality of every AS in V for a period of time T , we can 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 $\tilde{BC}_j(i)$ where $i \in V$ and $j \in T$.

In this paper, we calculate $\tilde{\mathbf{BC}}$ by analysing BGP data collected by RIPE and Route Views. 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 (every 8 hours in RIPE and every 2 hours in Route Views) and updates (every 5 minutes in RIPE and every 15 minutes in Route Views). Therefore, we can get routes from those peering ASes to any other ASes in the Internet, and compute variance of AS betweenness centrality by calculating changes in 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 [13]. We use this data to sample the Internet's behaviour.

2.2 Identify characteristics of rerouting

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 $\tilde{BC}_t = \langle \tilde{BC}_t(v_1), \tilde{BC}_t(v_2), \dots, \tilde{BC}_t(v_n) \rangle$, where $v_i \in V$. Higher magnitude of variation indicates more route changes. To characterize aggregated time of rerouting, 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.

Worst affected components The topological dimension of variance matrix includes all 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 $\tilde{BC}(v) = \langle \tilde{BC}_{t_1}(v), \tilde{BC}_{t_2}(v), \dots, \tilde{BC}_{t_m}(v) \rangle$, where $t_i \in T$. Higher diversity of this sequence of data indicates the AS experiences more unstable routing changes. Therefore, we use 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 descend instability.

A summation over a continuous time of elements in $\tilde{BC}(v)$ calculates changes of v 's betweenness centrality during that period of time. Formally speaking, $\tilde{BC}_{(t_p, t_{p+q})}(v)$ represents changes of v 's betweenness centrality between time t_p and t_{p+q} . It can be computed as $\tilde{BC}_{(t_p, t_{p+q})}(v) = BC_{t_{p+q}}(v) - BC_{t_p}(v) = BC_{t_{p+q}}(v) - BC_{t_{p+q-1}}(v) + BC_{t_{p+q-1}}(v) - BC_{t_{p+q-2}}(v) + \dots + BC_{t_{p+1}}(v) - BC_{t_p}(v) = \sum_{i=p+1}^{p+q} \tilde{BC}_{t_i}(v)$. 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 back to the origin state. These two aspects are considered to be the topological features of rerouting.

Correlative and backup paths Elements in networks are not independent. The effect of an event originated from one point will spread out over the network. 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 parts and backup paths of original routes. Synchronization of routing changes indicates common causes beneath them. It is of great significance for the following diagnose to identify and separate these ASes.

We develop a divisive hierarchical clustering algorithm [14] to analyse this synchronization. The basic idea is to evaluate the similarity between patterns of any two ASes' changes. As shown in Fig. 1, 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 value 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 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. However, 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 [8]. If the difference is marginal and less than a threshold θ , 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.

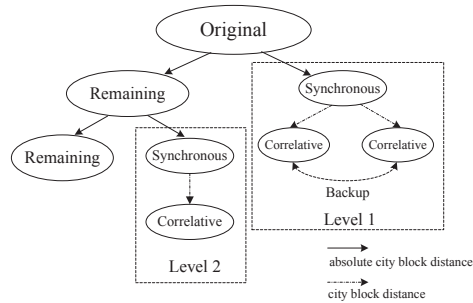


Fig. 1. Clustering algorithm for analysing synchronization of route changes

3 Rerouting after Japan Earthquake

After Japan's earthquake on Mar. 11th, news reported that some of trans-Pacific cable systems were damaged. Our network measurement system OneProbe observed regional path-quality degradation of the Internet during Mar. 11th and 12th. In this section, we apply our method to this event for better understanding the relationships among cable fault, inter-domain rerouting and end-to-end performance degradation.

We collected the data set from BGP routing tables and updates of RIPE and Route Views during Mar. 9th to 12th, then converted them into a sequence of 10-minute time slots so that the routing changes can be distinguished in fineness of

10 minutes. By extracting AS paths from RIBs in conjunction with updates, we can calculate every AS’s betweenness centrality in each time slot and construct a variance matrix.

We first limit the scope of our analysis to the routes from RRCs to HARNET’s IP prefixes monitored by OneProbe. Results associated with these specific destinations can help us infer the causes of reported performance degradation. After that, we further analyse the routes for the entire routed IP prefix space in order to understand the global routing state.

3.1 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. Further findings listed below show details of reroutings from different perspectives.

Fig. 2 shows the mean values of vector $|\tilde{BC}_t|$ wrt. time t , i.e., μ_t . 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 [17]. 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. 2. Aggregated time of route changes with HARNET

Fig. 3 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 . 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, transmitting data traffic for it. This result implies that traffic passing through these 3

ASes shifts most frequently. According to TeleGeography’s report [1], the Hong Kong-based cable-network operator PacNet reported damage 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 causes of changes on 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 unstable states, few of them have changed after the instability.




Fig. 3. Worst affected ASes with HARNET

We use the divisive hierarchical clustering algorithm described in Sec. 2 to exploit 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. 4. The variances decrease with increased levels. We stop at level 3 since changes are negligible after that. Fig. 5 shows the topology of AS paths from all RRCs to HARNET. Every end point is an AS in which a RRC is deployed. 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 evidence 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.

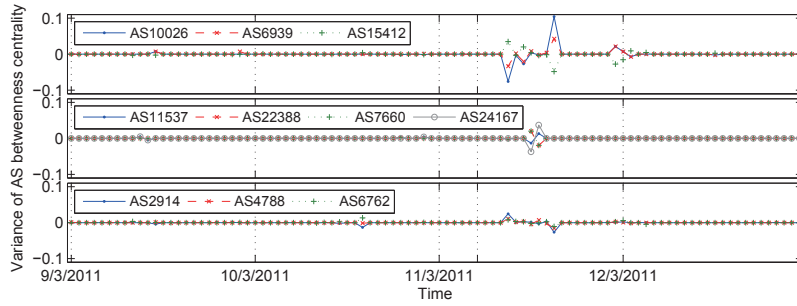


Fig. 4. Top 3 levels of synchronization with HARNET




Fig. 5. Topology of AS paths from all RRCs to HARNET

Moreover, in Fig. 4, positive and negative variances of an AS alternate over time, which indicates flapping routes passing through the AS. By flapping route here, we define it as the route changes from one AS path to another, then changes back to the original one again, e.g. from AS path $A \rightarrow B \rightarrow C$ to $A \rightarrow D \rightarrow E \rightarrow C$ then back to $A \rightarrow B \rightarrow C$. 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. Some routes don't just flap only once. The maximum time is 73, and 103 of them are more than 20 times. However, 96.9% of these routes are stick to the original one after flapping.

According to our results, we infer that the performance degradation on the paths from other places to HARNET in Hong Kong is mainly due to PacNet's cable fault 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.

3.2 Rerouting in the global Internet

In order to better understand inter-domain routing state of the entire Internet, we extend the destinations from HARNET's prefixes to all routed IP prefixes in BGP. 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.

Fig. 6 shows values of μ_t in 4 days. We can see that there are lots of spikes over time, probably due to BGP dynamics caused by numerous events in the Internet. But since earthquake and cable faults could result in persistent instability of routing system, we focus on the longest unusual changing time. We use the same method in previous section to identify the dominant period which lasts from 8:50 to 15:20 on 11th.




Fig. 6. Aggregated time of route changes in the Internet

Fig. 7 shows ranked ASes with their δ_v in a log-log scale, which indicates an inhomogeneous distribution of them. By separating them into 2 clusters, we get 17 worst affected ASes. They are AS10026, AS6939, AS174, AS3257, AS2914, AS6453, AS3549, AS3491, AS3356, AS7473, AS6762, AS3216, AS15412, AS6461, AS209, AS721 and AS27064, ordered by descend δ_v . Remarkably, PacNet (AS10026) is still top 1 in the entire Internet. It is evident that PacNet's regional cable faults have a vast impact on inter-domain rerouting of the Internet. Moreover, the changing ratio of AS10026's betweenness centrality during

the dominant period of rerouting is -21.91%. Unlike routes to Hong Kong, a large amount of routes which passing through PacNet and destined for other places shift away after unstable rerouting period.




Fig. 7. Worst affected ASes in the Internet

The synchronization of routing behaviours in the Internet is not as clear as that with HARNET. Fig. 8 is a demonstration of level 1 synchronization. Further analysis shows that the synchronization of routing behaviours is associated with destinations. It is more effective to study routings with topological-neighbour destinations.




Fig. 8. Level 1 synchronization in the Internet

4 Related work

Previous researches about inter-domain routing basically rely on analysing 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 [5]. Xu et al. employed principal component analysis (PCA) to transform and reduce BGP updates into different AS

clusters affected by distinct major events [6]. Zhang et al. proposed an instance-learning framework to identify anomalous BGP-route advertisements [7]. And Li et al. measured the impact of Internet earthquakes based on deviations from normal BGP update dynamics [8]. They employed BGP dynamics in different ways to achieve their different goals. However, it is a challenging work because BGP dynamics are inherent complex and full of outliers. In this paper, our intent is to diagnose network’s problem by examining changes of inter-domain routes. Hence we focus on changes of AS paths recorded in BGP routing tables and updates, which directly reflect inter-domain routing behaviours in AS level.

Fukuda et al., Cho et al. and Renesys reported their studies about limited impacts of Japan earthquake on the structure and routing dynamics of the regional Internet [15–17]. Especially, by examining BGP messages collected from local ISPs in Japan, Cho et al. claimed that almost nothing to see in BGP. Bischof et al. leveraged data collected by BitTorrent to identify specific regions and network links where Internet usage and connectivity were most affected after the earthquake [18]. Unlike the local routing information or specific application data that they use, we rely on publicly available control-plane data and focus on inter-domain routing of the Internet in a global scope. This method enables us to study the side effect of this earthquake on communications in Hong Kong and other places in the world.

5 Conclusion and future work

In this paper, we employ a social network-inspired metric, the variance of AS betweenness centrality, to analyse inter-domain routing behaviours after Japan earthquake. We infer that paths to Hong Kong are affected by unstable routing state caused by a cable fault after the earthquake. Besides, according to our extended analysis, the operator of this cable is also the worst impacted AS in the entire Internet, affecting paths to other destinations. We propose a method of using control-plane information to diagnose data-plane problems with publicly available BGP data. It is a useful and practical approach for understanding then improving the Internet functioning.

As an ongoing work, we will examine the routing states associated with other particular regions so as to achieve full spectrum vision of this event. By further exploring router-level routes presented in this work, we are able to correlate the paths with their geographical locations, then map them to submarine cables to better understand the relationship between logical routes and physical links. We will also apply our method to analysing other network events such as DDoS attacks, prefix hijackings, large-scale power outages and so forth.

Acknowledgement

This work is partially supported by 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; the National Natural

Science Foundation of China, under Grant No. 61003303, 61070199, 60903185; Program for Changjiang Scholars and Innovative Research Team in University (No. IRT1012); Program for Science and Technology Innovative Research Team in Higher Educational Institutions of Hunan Province (Network Technology, NUDT); and Hunan Province Natural Science Foundation of China (11JJ7003).

References

1. 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>.
2. Xiapu Luo, Edmond W. W. Chan and Rocky K. C. Chang. Design and Implementation of TCP Data Probes for Reliable and Metric-Rich Network Path Monitoring. In *Proc. USENIX Annual Tech. Conf.*, 2009.
3. RIPE RIS. <http://www.ripe.net/data-tools/stats/ris/routing-information-service>.
4. University of Oregon Route Views Project. <http://www.routeviews.org/>.
5. J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz. BGP routing dynamics revisited. In *ACM SIGCOMM Computer Communication Review*, 37(2):7-16, 2007.
6. 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.
7. J. Zhang, J. Rexford, and J. Feigenbaum. Learning-Based Anomaly Detection in BGP Updates. In *Proc. ACM SIGCOMM workshop on Mining network data*, 2005.
8. J. Li, and S. Brooks. I-seismograph: Observing and measuring Internet earthquakes. In *Proc. IEEE INFOCOM*, 2011.
9. L. C. Freeman. A Set of Measures of Centrality Based on Betweenness. In *Sociometry*, 40(1):35-41, 1977.
10. P. Holme, and B. J. Kim. Vertex overload breakdown in evolving networks. In *Phys. Rev. E* 65, 066109 (2002).
11. L. da F. Costa, F. A. Rodrigues, G. Travieso, and P. R. Villas Boas. Characterization of Complex Networks: A Survey of measurements. In *Advances in Physics*, 56(1):167-242, 2007.
12. L. Gao. On inferring autonomous system relationships in the internet. In *IEEE/ACM Transactions on Networking*, 9(6):733-745, 2001.
13. A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé. Analysis of Country-wide Internet Outages Caused by Censorship. In *Proc. ACM/USENIX IMC*, 2011.
14. J. Han and M. Kamber. *Data Mining: Concepts and Techniques, 2/e*. Morgan Kaufmann Publishers, 2006.
15. K. Fukuda, M. Aoki, S. Abe, Y. Ji, M. Koibuchi, M. Nakamura, S. Yamada, and S. Urushidani. Impact of Tohoku Earthquake on R&E Network in Japan. In *Proc. ACM SWID*, 2011.
16. 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.
17. Renesys blog. Japan Quake. <http://www.renesys.com/blog/2011/03/japan-quake.shtml>
18. Z. S. Bischof, J. S. Otto, and F. E. Bustamante. Distributed Systems and Natural Disasters. In *Proc. ACM SWID*, 2011.