Characterizing Large-scale Routing Anomalies: A Case Study of the China Telecom Incident

Full version from December 13, 2012

Rahul Hiran[1], Niklas Carlsson[1], and Phillipa Gill[2]⋆

[1] Linköping University, Sweden
[2] Citizen Lab, Munk School of Global Affairs
University of Toronto, Canada

Abstract. China Telecom's hijack of approximately 50,000 IP prefixes in April 2010 highlights the potential for traffic interception on the Internet. Indeed, the sensitive nature of the hijacked prefixes, including US government agencies, garnered a great deal of attention and highlights the importance of being able to characterize such incidents after they occur. We use the China Telecom incident as a case study, to understand (1) what can be learned about large-scale routing anomalies using public data sets, and (2) what types of data should be collected to diagnose routing anomalies in the future. We develop a methodology for inferring which prefixes may be impacted by traffic interception using only control-plane data and validate our technique using data-plane traces. The key findings of our study of the China Telecom incident are: (1) The geographic distribution of announced prefixes is similar to the global distribution with a tendency towards prefixes registered in the Asia-Pacific region, (2) there is little evidence for subprefix hijacking which supports the hypothesis that this incident was likely a leak of existing routes, and (3) by preferring customer routes, providers inadvertently enabled interception of their customer's traffic.

Keywords: Measurement, Routing, Security, Border Gateway Protocol

1 Introduction

On April 8, 2010, AS 23724, an autonomous system (AS) owned by China Telecom, announced approximately 50,000 prefixes registered to other ASes. These prefixes included IPs registered to the US Department of Defense [8], which caught the attention of the US-China Economic and Security Review Commission [5]. Unlike previous routing misconfigurations [6, 17], China Telecom's network had the capacity to support the additional traffic attracted [4]. Further, there is ample data-plane evidence suggesting that during the incident, Internet traffic was reaching its correct destination. This unique situation is what led some to suggest this was an attempt to intercept Internet traffic.

⋆ Data sets available at: http://www.ida.liu.se/∼nikca/papers/pam13.html

While the China Telecom incident has garnered attention in blogs [4,8], news outlets [16], and government reports [5], there has been no academic attempt to understand this incident. This dearth of understanding is especially apparent when considering the many questions that remain unanswered about this incident. These include (1) understanding properties of the hijacked prefixes, (2) quantifying the impact of the event in terms of subprefix hijacking, and (3) explaining how interception was possible. We tackle these questions using publicly available control- and data-plane measurements and highlight what types of data would be useful to better understand routing anomalies in the future. We emphasize that while we are able to characterize the incident and show evidence supporting the hypothesis that this incident appears to be an accident, there is currently no way to distinguish between "fat finger" incidents and those that have malicious intent based on empirical data alone.

1.1 Insecurity of the Internet's Routing System

Routing security incidents have happened repeatedly over the past 15 years [6,12, 17]. These incidents involve an AS originating an IP prefix without permission of the autonomous system (AS) to which the prefix is allocated: *hijacking*. Usually when hijacks happen, the misconfigured network either does not have sufficient capacity to handle the traffic [17] or does not have an alternate path to the destination [6]. In these cases, the impact of the incident is immediately felt as a service outage or interruption of connectivity.

More troubling, are cases of traffic *interception*, where traffic is able to flow through the hijacking AS and on to the intended destination. Without continuous monitoring of network delays or AS paths [3], incidents such as these are difficult to detect, thus creating opportunity for the hijacker to monitor or alter intercepted traffic. Traffic interception was demonstrated in 2008 [20] and more recently occurred during the China Telecom incident [8].

Since the China Telecom incident involved interception, measuring its impact is extremely difficult without extensive monitoring infrastructure. We define criteria that allow us to infer potential interceptions using only control-plane data [18]. We use data-plane measurements [14] to validate our criteria and characterize the AS topologies that allowed for inadvertent interception.

1.2 Key Insights

The geographic distribution of announced prefixes does not support targeted hijacking. The distribution of announced prefixes is similar to the geographic distribution of all globally routable prefixes with a tendency towards prefixes in the Asia-Pacific region.

The prefixes announced match existing routable prefixes. We observe that $> 99\%$ of the announced prefixes match those existing at the Routeviews monitors. This supports the conclusion that the announced prefixes were a subset of AS 23724's routing table.

Providers inadvertently aided in the interception of their customers' traffic. Many networks that routed traffic from China Telecom to the correct destination did so because the destination was reachable via a customer path which was preferred over the path through China Telecom (a peer).

2 Related Work

While China Telecom incident occurred in April 2010, it received little attention [4,16] until November 2010 when the US-China Economic and Security Review Commission published their report to congress [5] which included a description of the event. After the release of the report, the incident received attention in news articles and was investigated by some technically-oriented blogs [8,13].

BGPMon, an organization that provides monitoring and analysis of BGP data, performed the first investigation of the China Telecom incident [4]. Using control-plane measurements of BGP messages, they were able to identify anomalous updates as those where the path terminated in "4134 23724 23724." They also study the geographic distribution of the hijacked prefixes and find that the majority of hijacked prefixes belong to organizations in the US and China.

Using both control- and data-plane data, Renesys confirmed the geographic distribution of hijacked prefixes observed by BGPMon [8]. Using traceroute, Renesys was also able to show that network traffic was able to pass into China Telecom's network and back out to the intended destination. Further analysis was performed by Arbor Networks [13] which focused on understanding how much traffic was diverted into China Telecom using traffic flows observed through ASes participating in the ATLAS project [2]. They do not observe a significant increase in traffic entering AS 4134 on the day of the incident.

In contrast to the blog entries, our focus is on analyzing the incident using only publicly available data to understand what can be learned using today's public data and what types of data should be collected in the future.

3 Methodology

To characterize the events that took place on April 8, 2010, we use a combination of publicly available control- and data-plane measurements [7, 14, 18]. In this section, we provide background on China Telecom's network and the data sets used in our analysis.

3.1 China Telecom's Network

China Telecom is the 11th largest ISP on the Internet [8] and maintains multiple ASes partitioning their resources into different geographic regions (*e.g.,* provinces) and types (*e.g.,* data centers vs. regional networks). Many of these ASes are found as customers of AS 4134 in the AS-graph [7] and can be further identified using `whois` data. Indeed, the erroneous BGP updates originate from AS 23724 which is actually an AS owned by China Telecom and is located in Beijing.

Table 1. Summary of control-plane updates matching the attack signature.

Monitor (Location)	Number of Updates	Number of Unique Prefixes
LINX (London, England)	60,221	11,413
DIXIE (Tokyo, Japan)	80,175	15,773
ISC PAIX (Palo Alto, CA)	123,723	35,957
Route-vews2 (Eugene, OR)	216,196	29,998
Route-views4 (Eugene, OR)	49,290	18,624
Equinix (Ashburn, VA)	44,793	13,250
BGPMon list	-	37,213
Total	574,398	43,357

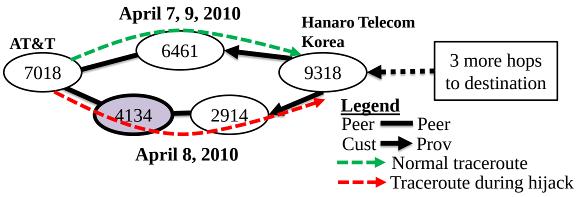

Fig. 1. Interception observed in the traceroute from planet2.pittsburgh.intel-research.net to 125.246.217.1 (DACOM-PUBNETPLUS, KR).

3.2 Control-plane Measurements

We use a combination of BGP updates [18] and topology data [7] to characterize the China Telecom incident.

BGP updates. We use Routeviews monitors as a source of BGP updates from around the time of the attack. We consider updates with the path attribute ending in "4134 23724 23724" as belonging to the incident [4].[3] Table 1 summarizes the updates and prefixes matching this signature from the Routeviews monitors.

Topology data. We use the Cyclops AS-graph from April 8, 2010 [7] to infer the set of neighbors of China Telecom and their associated business relationships. Knowing the neighbors of China Telecom is particularly important when identifying ASes that potentially forwarded traffic in (and out of) China Telecom during the incident.

3.3 Data-plane Measurements

We use data-plane measurements from the iPlane project [14] and extract traceroutes transiting China Telecom's network on April 8, 2010. We first map each IP in the traceroute to the AS originating the closest covering prefix at the time of the traceroute. If we observe a traceroute AS-path that does not contain China Telecom (AS 4134 or AS 23724) on April 7 or 9, that *does* contain these networks

[3] All but 36 prefixes originated by AS 23724 match this signature.

on April 8, we conclude that this traceroute was impacted by the incident. Further, if we observe a traceroute that was impacted, and the final AS in the path is not AS 4134 or 23724, we conclude that this traceroute was intercepted. Figure 1 shows a traceroute where interception was observed. This traceroute only transits AS 4134 (China Telecom) on April 8 and is able to reach the destination through AS 2914 (NTT) who provides transit for AS 9318 (Hanaro Telecom). In total, we observed 1,575 traceroutes transiting China Telecom on April 8. Of these, 1,124 were impacted by the routing incident and 479 were potentially intercepted, with 357 of these receiving a successful response from the target.

3.4 Limitations

We face limitations in existing data sets as we reuse them for the unintended task of analyzing a large-scale routing anomaly.

Inaccuracies in the AS-graph. AS-graphs suffer from inaccuracies inferring AS-relationships (*e.g.,* because of Internet eXchange Points (IXPs) [1]) and poor visibility into peering links [19]. These inaccuracies impact our analysis of interception which uses the AS-graph to infer China Telecom's existing path to a destination. We discuss this limitation in more detail in Section 5.3.

Inaccurate IP to AS mappings. We note that our mapping of IP addresses to ASes may be impacted by IXPs or sibling ASes managed by the same institution [15]. Since our primary concern is paths that enter China Telecom *only* on April 8 the impact of siblings (*e.g.,* per-province ASes managed by China Telecom) should be mitigated. This is because paths to China Telecom's siblings would normally transit China Telecom's backbone AS 4134.

4 Impact of the China Telecom Hijack

We now consider the impact of the China Telecom incident in terms of the prefixes that were announced.

4.1 What is the geographic distribution of the announced prefixes?

Figure 2 shows a breakdown of the prefixes that were hijacked by country, with and without excluding prefixes owned by AS 4134. We see the bulk of prefixes are registered to organizations in US and China, followed by Korea, Australia and Mexico, which is consistent with observations made by BGPMon [4].

Was it random? Figure 2 also plots the geographic distribution of all routable prefixes on the Internet. Here we can see a disproportionate number of Chinese prefixes (especially belonging to AS 4134) being hijacked. Additionally, when comparing hijacked prefixes to the global distribution of prefixes there appears to be little evidence for attack. Indeed, the US shows fewer prefixes being hijacked than would be expected based on the global distribution, while countries in the Asia-Pacific region (e.g., China, Korea, Australia) have more hijacked prefixes.

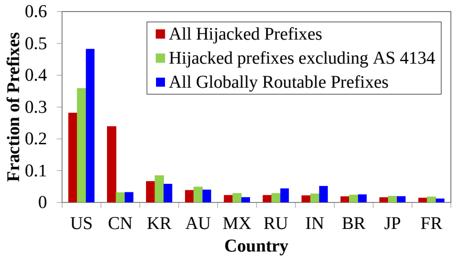

Fig. 2. Top 10 countries impacted by the China Telecom incident.

Table 2. Organizations most impacted by the China Telecom incident.

	Overall		Subprefix Hijacks	
Rank	Prefixes	Organization	Prefixes	Organization
1	9,296	China Telecom (AS 4134)	8,614	China Telecom (AS 4134)
2	1,573	Time Warner (AS 4323)	371	China Educ/Research (AS 4538)
3	1,229	Korea Telecom (AS 4766)	11	China Telecom (AS 38283)
4	739	AT&T (AS 7018)	9	Telecom Holding (AS 34590)
5	569	Tata (AS 4755)	4	Cisco Systems (AS 109)

4.2 Which organizations were most impacted?

Table 2 considers the organizations most impacted by the China Telecom incident; both overall, and when only considering subprefix hijacks. Organizations with the most prefixes announced tend to be peers of AS 4134. Indeed, direct neighbors of China Telecom are most adversely impacted with an average of 85 prefixes hijacked vs. 9 prefixes hijacked for all impacted ASes.

Critical networks were subject to hijacking. While they do not make the top five list, China Telecom announced some critical US prefixes. Government agencies such as Department of Defense, United States Patent and Trademark Office, and Department of Transport were impacted. In addition, commercial entities such as Apple, Cisco, DE Shaw, HP, Symantec and Yahoo! were impacted. This fact was also noted in the report made to the US congress [5].

4.3 Were any of the announcements subprefix hijacks?

We now consider the length of the prefixes announced by China Telecom relative to existing routes. If the event was simply a leak of routes contained in the routing table, it should be the case that China Telecom's prefixes will be the same length as existing routes. Additionally, prefix length can shed light on the impact of the incident since more specific prefixes are preferred. For each of the six Routeviews monitors (Section 3.2), we use the RIB tables as seen on April 7 to derive the existing prefix lengths. Route aggregation means that the prefix length observed varies between the vantage points.

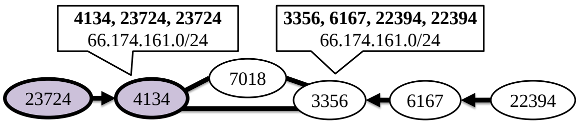

Fig. 3. Example topology that allows for interception of traffic.

Subprefix hijacking was extremely rare. In total, 21% (9,082) of the prefixes were longer than existing prefixes at all six monitors. However, 95% (8,614) of these prefixes belong to China Telecom (Table 2). Most of the observed subprefix hijacking is due to poor visibility of Chinese networks (AS 4134, 4538, and 38283) at the monitors. Excluding these networks, we observe < 1% (86) prefixes being subprefix hijacked. The lack of subprefix hijacks supports the conclusion that the incident was caused by a routing table leak.

Our analysis highlights the importance of using multiple vantage points to mitigate the impact of route aggregation on results. Indeed, an additional 236 prefixes were subprefix hijacked at at least one vantage point.

5 The Mechanics of Interception

The fact that traffic was able to flow through China Telecom's network and onto the destination is highly unusual. We now discuss how interception may occur accidentally, based on routing policies employed by networks.

5.1 How was interception possible?

Two key decisions, when combined with inconsistent state within China Telecom's network, allow for traffic to be intercepted. These properties have also been discussed in related work [3]. We illustrate them with an example from the China Telecom incident (Figure 3). This figure was derived using a combination of BGP updates [18] and a traceroute observed during the incident [8].

Decision 1: AT&T (AS 7018) chooses to route to China Telecom. In Figure 3, AT&T (AS 7018) has two available paths to the prefix. However, since the path advertised by China Telecom (AS 4134) is shorter, AT&T (AS 7018) chooses to route to China Telecom.

Decision 2: Level 3 (AS 3356) chooses *not* to route to China Telecom. In order for traffic to leave China Telecom's network and flow on to the intended destination, China Telecom requires a neighbor that *does not* choose the path it advertises. In the example above, this occurs when Level 3 (AS 3356) chooses to route through its customer Verizon (AS 6167) instead of through its peer China Telecom (AS 4134). Thus, China Telecom can send traffic towards Level 3 and have it arrive at the intended destination.

We next characterize what causes these decisions to be made using a combination of control- and data-plane data.

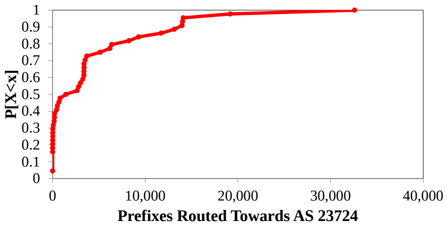

Fig. 4. Cumulative distribution function of prefixes each neighboring AS routed towards China Telecom.

Table 3. Neighbors that routed the most prefixes to China Telecom.

Rank	# of Prefixes	% of Prefixes	Organization
1	32,599	75%	Australian Acad./Res. Net. (AARNet) (AS 7575)
2	19,171	44%	Hurricane Electric (AS 6939)
3	14,101	33%	NTT (AS 2914)
4	14,025	32%	National LambdaRail (AS 11164)
5	13,970	32%	Deutsche Telekom (AS 3320)

5.2 How many ISPs chose to route to China Telecom?

We first consider how many ISPs made **Decision 1**. We observe 44 ASes routing traffic towards China Telecom, with each AS selecting the path through China Telecom for an average of 4,342 prefixes. Figure 4 shows the cumulative density function of the number of prefixes each AS routes to China Telecom. The distribution of prefixes each AS routes to China Telecom is highly skewed, with some ASes being significantly more impacted than others. The top five ASes are summarized in Table 3, which highlights the role of geography in the hijack, with networks operating in Europe and Asia-Pacific regions being most impacted. Academic networks (AARNet and National LambdaRail) are also heavily impacted.

We consider the paths these five networks have to the victim prefixes and find that 98% have a path length of three or more using the standard routing policy model [9] (96% if we assume shortest path routing). This means that the three-hop path ("4134 23724 23724") announced by China Telecom would have been shorter than their existing path.

5.3 Which prefixes were intercepted?

We develop a methodology to locate potentially intercepted prefixes using control-plane data. Control-plane data has the advantage that it may be passively collected in a scalable manner. We validate our technique and further analyze the interception that occurred using data-plane measurements [14] (Section 5.4).

We use the following methodology to locate potentially intercepted prefixes using only control-plane data. First, for each hijacked prefix, we use the Cyclops AS-graph (discussed in Section 3.2) and a standard model of routing policies [9],

Table 4. Organizations with the most potentially intercepted prefixes.

Rank	Prefixes	Organization
1	712	AT&T (AS 7018)
2	174	Eli Lilly (AS 4249)
3	136	China Telecom (AS 4134)
4	120	Sprint (AS 1239)
5	108	Level 3 (AS 3356)

to compute China Telecom's best path to the prefix.[4] Next, for each of these paths, we check whether the next-hop on China Telecom's best path to the destination was observed routing to China Telecom for the given prefix.

We observe 68% of the hijacked prefixes potentially being intercepted; however, 85% of these prefixes are observed being intercepted via AS 9304, a customer of China Telecom, which may be an artifact of poor visibility of the Routeviews monitors. Excluding paths through AS 9304, we observe a total of 10% (4,430) prefixes potentially being intercepted. Table 4 summarizes the organizations with the most intercepted prefixes. In the case of AT&T, Sprint and Level 3 these ASes are providers and peers to China Telecom that still provided China Telecom (a peer) with paths to their prefixes. Additionally, some Department of Defense prefixes may have been intercepted as China Telecom potentially still had a path through Verizon.

Limitations. Our method is limited in two key ways. First, we may not observe all announcements made by China Telecom's neighbors (*i.e.,* we may incorrectly infer that they are *not* routing to China Telecom because their announcement is not seen by the Routeviews monitors). Second, we do not know which neighbor China Telecom would normally use to transit traffic for a given prefix and thus we have to infer it based on topology measurements and a routing policy model.

Validation. Without ground-truth data it is difficult to quantify the inaccuracies of our methodology as we may both over- or under-estimate the potential for interception. We use the data-plane measurements described in Section 3.3 to validate our methodology. Of the 479 traceroutes that were intercepted, 319 (66%) were observed in prefixes detected as intercepted using our criteria. This inaccuracy stems from a lack of control-plane data which leads to the limitations mentioned above. With more complete data, our method could better identify potential interceptions.

5.4 Why neighboring ASes did *not* route to China Telecom?

We use data-plane measurements to understand why neighboring ASes chose not to route to China Telecom (**Decision 2**). In Figures 3 and 1, the reason that the neighboring AS does not route to China Telecom is because they have

[4] Since China Telecom does not normally transit traffic for the hijacked prefixes, we were unable to extract the paths normally used by China Telecom from Routeviews.

Table 5. Why networks chose not to route to China Telecom.

Reason	# of traceroutes	% of traceroutes
Had a customer path	139	39%
Had a shorter path	193	54%
Had an equally good path	18	5%
Other	7	2%

a path through a customer to the destination. However, this is only one reason an AS would choose not to route to China Telecom. We consider the cases of interception observed in the iPlane data and determine why the neighboring AS did not route to China Telecom using the Gao-Rexford routing policy model [9].

Table 5 summarizes the reasons neighbors of China Telecom did not route to China Telecom. Here we only consider the 357 traceroutes where interception was observed and a response was received from the target. The majority of neighbors handling intercepted traffic did not choose the China Telecom route because it was longer than their existing route for the prefix in question.

Providers inadvertently allowed interception of customer traffic. A significant fraction (39%) of neighbor ASes do not route to China Telecom because they have a path to the destination via a customer, such as AS 3356 in Figure 3. These providers inadvertently aided in the interception of their customer's traffic by forwarding China Telecom's traffic to the destination. While providers cannot control the traffic sent to them by neighboring ASes, it may be beneficial to monitor the neighbors sending traffic towards their customers for anomalies, so that customers may be alerted to potential interception events.

We observe seven traceroutes where it is unclear why the China Telecom path was not chosen. These traceroutes involved a provider to China Telecom who chose to route towards other customers likely the result of traffic engineering or static routes being used for the customer ASes.

6 Discussion

Using publicly available data sources we have characterized the China Telecom incident that occurred in April 2010. Our study sheds light on properties of the prefixes announced, and supports the conclusion that the incident was a leak of random prefixes in the routing table, but does not rule out malicious intent.

On diagnosing routing incidents. Our work highlights the challenge of understanding large-scale routing incidents from a purely technical perspective. While empirical analysis can provide evidence to support or refute hypotheses about root cause, it cannot prove the intent behind the incident. However, empirical analysis can provide a starting point for discussions about the incident.

On the available data. When the results of analysis can lead to real-world reaction it is important that the data used for analysis is as complete as possible and robustness/limitation of results are clearly stated. These two properties can be achieved by increasing the number of BGP monitors [11] and performing careful analysis of robustness and limitations [10].

Acknowledgments

The authors thank the reviewers and our shepherd Olaf Maennel for constructive suggestions, that helped improve the paper. This work benefited from discussions with Jennifer Rexford and Andy Ogielski. We thank Monia Ghobadi, Sharon Goldberg, and Jennifer Rexford for comments on drafts. We thank Doug Madory for assistance with the analysis in Section 3.3.

References

1. B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a large European IXP. In *Proc. of ACM SIGCOMM*, 2012.
2. ATLAS - Arbor Networks, 2012. http://atlas.arbor.net.
3. H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the Internet. In *Proc. of ACM SIGCOMM*, 2007.
4. BGPMon. China telecom hijack, 2010. http://bgpmon.net/blog/?p=282.
5. D. Blumenthal, P. Brookes, R. Cleveland, J. Fiedler, P. Mulloy, W. Reinsch, D. Shea, P. Videnieks, M. Wessel, and L. Wortzel. Report to Congress of the US-China Economic and Security Review Commission, 2010. http://www.uscc.gov/annual_report/2010/annual_report_full_10.pdf.
6. M. Brown. Renesys blog: Pakistan hijacks YouTube. http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml.
7. Y. Chi, R. Oliveira, and L. Zhang. Cyclops: The Internet AS-level observatory. *ACM SIGCOMM Computer Communication Review*, 2008.
8. J. Cowie. Renesys blog: China's 18-minute mystery. http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml.
9. L. Gao and J. Rexford. Stable Internet routing without global coordination. *Transactions on Networking*, 2001.
10. P. Gill, M. Schapira, and S. Goldberg. Modeling on quicksand: Dealing with the scarcity of ground truth in interdomain routing data. *ACM SIGCOMM Computer Communication Review*, 2012.
11. E. Gregori, A. Improta, L. Lenzini, L. Rossi, and L. Sani. On the incompleteness of the AS-level graph: a novel methodology for BGP route collector placement. In *ACM Internet Measurement Conference*, 2012.
12. V. Khare, Q. Ju, and B. Zhang. Concurrent prefix hijacks: Occurrence and impacts. In *ACM Internet Measurement Conference*, 2012.
13. C. Labovitz. China hijacks 15% of Internet traffic?, 2010. http://ddos.arbornetworks.com/2010/11/china-hijacks-15-of-internet-traffic/.
14. H. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An information plane for distributed services. In *In Proc. of OSDI*, 2006.
15. Z. Mao, J. Rexford, J.Wang, and R. H. Katz. Towards an accurate AS-level traceroute tool. In *Proc. of ACM SIGCOMM*, 2003.

16. R. McMillan. A Chinese ISP momentarily hijacks the Internet, 2010. http://www.nytimes.com/external/idg/2010/04/08/08idg-a-chinese-isp-momentarily-hijacks-the-internet-33717.html.
17. S. Misel. “Wow, AS7007!”. Merit NANOG Archive, 1997. http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html.
18. U. of Oregon. Route views project. http://www.routeviews.org/.
19. R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. Quantifying the completeness of the observed internet AS-level structure. *UCLA Computer Science Department - Techical Report TR-080026-2008*, Sept 2008.
20. A. Pilosov and T. Kapela. Stealing the Internet: An Internet-scale man in the middle attack, 2008. Presentation at DefCon 16, http://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-kapela.pdf.