Impact of Prefix Hijacking on Payments of Providers

Pradeep Bangera, Sergey Gorinsky
Institute IMDEA Networks
(Madrid Institute for Advanced Studies in Networks)
Avenida del Mar Mediterraneo, 22
Leganes, Madrid, 28918, Spain
Email: pradeep.bangera@imdea.org, sergey.gorinsky@imdea.org

***Abstract*—Whereas prefix hijacking is usually examined from security perspectives, this paper looks at it from a novel economic angle. Our study stems from an observation that a transit AS (Autonomous System) has a financial interest in attracting extra traffic to the links with its customers. Based on real data about the actual hijacking incident in the Internet, we conduct simulations in the real AS-level Internet topology with synthetic demands for the hijacked traffic. Then, we measure traffic on all inter-AS links and compute the payments of all providers. The analysis of our results from technical, business, and legal viewpoints suggests that hijacking-based traffic attraction is a viable strategy that can create a fertile ground for tussles between providers. In particular, giant top-tier providers appear to have the strongest financial incentives to hijack popular prefixes and then deliver the intercepted traffic to the proper destinations. We also discuss directions for future research in the area of hijacking-based traffic attraction.**

I. Introduction

Thousands of ISPs (Internet Service Providers) contribute their communication resources to provide universal connectivity to more than one billion Internet users. The communication infrastructure of an ISP consists of a single or multiple ASes (Autonomous Systems), with some large ISPs – such as AT&T – owning more than 10 ASes. The AS-level topology of the Internet is hierarchical in its essence. A vast majority of ISPs are relatively small and specialize in providing Internet access to end users. A much smaller family of large providers form the Internet core and deliver traffic between other ISPs.

The creation and maintenance of the ISP communication infrastructure involve substantial costs. It is common for ISPs to operate as commercial entities and recover the costs by charging customers for consumed services. For example, a large provider that connects a small ISP to the rest of the Internet charges this small customer for the traffic transited in both directions of the communication link between the customer and provider. While such transit business relationships between ASes are frequent, some ISPs act toward each other as peers and exchange traffic over interconnecting links without any financial requital. The peering arrangement is typical for ISPs of a similar stature and is limited to balanced exchanges of traffic between their own pools of customers. For instance, peering can be an economically attractive option for two small ISPs in the same geographic area because maintenance of the local peering link is less costly than transit of the exchanged traffic through their shared upstream provider. Finally, a link that connects two ASes belonging to the same ISP is called a sibling link. Traffic on sibling links does not result in any inter-provider payment.

To deliver traffic in accordance with own economic preferences, ISPs utilize BGP (Border Gateway Protocol) [1] and IP (Internet Protocol) [2], the technical means for routing and forwarding in the Internet. The final destination of communicated data is identified by an IP address. An IP prefix is a succinct representation for contiguous IP addresses. Because IP addresses are assigned to an ISP in contiguous chunks, IP prefixes form a basis for scalable IP forwarding: an IP router quickly forwards an incoming IP datagram to the output link determined through the longest-prefix match in the forwarding table that compactly maps IP prefixes to output links [3]. The forwarding tables are constructed by routing protocols. In particular, BGP serves as a de facto standard protocol for routing between ASes. BGP is a path vector protocol where ASes send their neighbors announcements about AS-level paths to IP prefixes. The announcing AS either owns the advertised prefix or learns about the path to the prefix from another neighbor. The announced path lists the announcing AS as the first hop to the prefix. Based on various policies including the aforementioned economic considerations, each AS decides whether to use the learned paths for forwarding own traffic and to which neighboring ASes the path knowledge should be propagated.

Although BGP is a sophisticated protocol, it does not provide an ISP with a reliable mechanism to validate the path information announced by neighboring ISPs. IP prefix hijacking [4]–[12] is a general term for announcement practices that deviate from the expected BGP behavior and thereby divert traffic to different paths. Specific forms of prefix hijacking include pretending to own an IP prefix that belongs to another AS, announcing a shorter path to the prefix than the path announced by the prefix owner, and advertising a more specific prefix than the one announced by the original AS. Internet security experts have recognized the vulnerability of BGP to prefix hijacking and studied blackholing [13], eavesdropping [14], phishing [15], and other hijacking-based attacks. Various solutions such as S-BGP (Secure BGP) [16], announcement filtering [17], and pre-

fix registration databases [18] have been proposed but their effectiveness is limited in practice, e.g., because of serious challenges with deploying them broadly in the multi-provider Internet.

As the name suggests, prefix hijacking is usually viewed from security perspectives. In this paper, we investigate prefix hijacking from a novel economic angle. Because ISPs engage in transit business relationships, and make and receive payments for traffic communicated over customer-provider links, a provider has a direct financial interest in attracting extra traffic to its links with customers, e.g., when competing with another provider for transit services supplied to a multihomed customer [5], [19]. In comparison to the well-known technique of hot-potato routing [20] where a provider saves on its internal communication resources by handing over transit traffic to another ISP at the nearest interconnection, prefix hijacking enables more intricate economic tussles between ISPs, e.g., multi-hop competition for transit traffic from faraway ASes.

This paper studies the impact of prefix hijacking on inter-ISP relationships and, in particular, payments of providers. To assess whether prefix hijacking is a viable strategy for revenue boosting via traffic attraction, our research method uses real data as a basis for simulating hypothetical scenarios. In addition to technical aspects of hijacking-based traffic attraction, we discuss its feasibility from business, legal, and other pertinent perspectives.

The rest of this paper is organized as follows. Section II presents the general methodology of our studies. Based on data about a real prefix-hijacking incident in the Internet, Section III conducts simulations to evaluate the impact of prefix hijacking on BGP paths, inter-AS link traffic, and provider payments. Section IV assesses viability of prefix hijacking for revenue-boosting traffic attraction. Section V outlines directions for future research in this area. Section VI discusses related work. Section VII concludes the paper with a summary of our contributions.

II. Methodology

Confidentiality concerns limit public information about traffic, routing, and pricing inside the Internet. To deal with this challenge, our study combines real data with simulations and speculative reasoning.

The specific basis for our investigation is a real incident of prefix hijacking in the Internet: the hijacking of a YouTube prefix by Pakistan Telecom in February 2008 [21]. The incident has attracted significant attention and been actively discussed, e.g., on the NANOG (North American Network Operator Group) mailing list [22]. Besides, the Internet incorporates BGP announcement monitoring systems such as PHAS (Prefix Hijack Alert System) [7], RIPE RIS (Réseaux IP Européens Routing Information Service) [23] and BGPmon (BGP monitoring and analyzer tool) [24]. In our analysis, we rely on actual announcement data collected by the monitoring systems during the hijacking incident.

The adopted research method supplements the real-data analysis with simulations for two reasons. First, the available real data do not paint the full picture of the hijacking incidents. Second, simulations enable us to examine suppositional scenarios that are more suitable for revenue-boosting traffic attraction. Our choice for the simulation platform is C-BGP [25], a widely used simulator for BGP routing problems. While C-BGP is an advanced tool, it deviates from reality in some pertinent aspects. For example, the default valley-free configuration [26] of C-BGP is reasonable in general but occasionally yields different paths than those produced by actual BGP import and export policies, i.e., ISP rules for adopting and propagating BGP paths. Also, the C-BGP support for prefix hijacking is not comprehensive, e.g., the simulator does not include a feature where a transit AS poses as an intermediary and announces a more specific prefix than the one advertised by the original AS. This feature is relevant to our research because we are interested in traffic-attraction scenarios where the transit AS forwards the intercepted traffic to the original AS. We integrated the feature into C-BGP and supplied the modified version to the C-BGP developers [27].

To initialize C-BGP with a realistic contemporary AS-level topology of the Internet, we use a data set collected by CAIDA (Cooperative Association for Internet Data Analysis) [28] at the time of the real hijacking incident. The CAIDA data set classifies relationships between a pair of ASes as customer-provider (encoded as -1 in the data set), provider-customer (encoded as 1), and peering (encoded as 0). We remove from each data set all customer-provider pairs because of their redundancy: for every provider-customer relationship provided to C-BGP, the simulator automatically configures a transit link associated with both provider-customer and corresponding customer-provider relationships. Since C-BGP does not recognize sibling relationships between ASes, we also substitute sibling relationships (encoded as 2) with peering relationships. The number of the sibling relationships in the CAIDA data set is small, and the substitution has a negligible impact on the fidelity of the simulations.

Our simulations rely on synthetic traffic to evaluate the impact of prefix hijacking on inter-provider links. With each AS in the Internet-scale topologies, we associate a traffic demand for the advertised prefix. Then, we utilize C-BGP to determine the rate of traffic flowing in both directions of each inter-ISP link. We refer to this traffic rate as a *link load* of the inter-provider link.

Finally, we translate the link loads of inter-provider links into payments between ISPs. In doing this, we distinguish between transit and peering links and adopt the respective pricing functions from [29]–[31]. For a transit link, monthly payment is passed from the customer to the provider and calculated in $ (USA dollars) as

$$p_t = m_t \cdot v^{0.75} \tag{1}$$

where v represents the total traffic in both directions of the link in Kbps, and coefficient $m_t = 0.0675$ is such that 1 Mbps is priced at $12 [31]. The sublinear dependence on v reflects the economies-of-scale effect of paying less for a traffic unit as the traffic volume increases. Peering links rarely

incur charges that are directly connected to the exchanged traffic. Instead, the peering ASes share the costs of maintaining the connection, e.g., by paying an annual fee to an IXP (Internet eXchange Point) [31]. These costs are significantly smaller than transit charges for similar traffic settings but are not negligible nevertheless. Hence, following the guidelines from [30], [31], each of the two ASes in a peering relationship is assumed to pay a third party the same monthly amount of

$$p_e = m_e \cdot v^{0.4} \tag{2}$$

where v is raised to the smaller power of 0.4, and coefficient $m_e = 0.0631$ ensures that 1 Mbps is priced at \$1. Even though Equations 1 and 2 offer only rough estimates of actual prices, which vary with time and geographic region, the above pricing functions allow us to derive preliminary insights into the economic viability of revenue-boosting traffic attraction via prefix hijacking. For each AS, we partition all inter-provider links of the AS into three sets: set R contains the transit links where the AS acts as a provider, set C is for the transit links that involve the AS as a customer, and set E captures all peering links of the AS. Then, we compute overall monthly payment p of the AS as

$$p = \sum_{t \in R} p_t - \sum_{t \in C} p_t - \sum_{e \in E} p_e \tag{3}$$

with positive values denoting financial gains, and negative values representing financial losses.

III. YouTube hijacking by Pakistan Telecom

YouTube owns AS 36561 with five assigned prefix spaces according to the RIPE RIS Dashboard [23]. 208.65.152.0/22 represents one such space and is the prefix that attracts a majority of YouTube-addressed traffic. On the 24th of February in 2008, AS 17557 belonging to PTCL (Pakistan Telecommunication Company Limited) hijacked YouTube traffic for approximately two hours and fourteen minutes by announcing the more specific prefix 208.65.153.0/24. The intention of the hijacking was to block access to YouTube within the state of Pakistan but the impact was significantly more far-reaching because PTCL announced 208.65.153.0/24 also to its provider PCCW Global (AS 3491), and the latter advertised globally the bogus PTCL paths for the longer prefix. Consequently, PTCL became a black hole that attracted and discarded packets sent to YouTube from all over the global Internet. YouTube detected the sharp decrease in its incoming traffic and reacted by announcing the even more specific prefix 208.65.153.0/25. The countermeasure restored some traffic flow to YouTube, yet PTCL remained able to attract a nontrivial fraction of YouTube-addressed traffic due to the path length and other factors that affect the routing policies of various ASes [21].

In accordance with our general simulation methodology from Section II, we initialize C-BGP with an AS-level Internet topology as per the CAIDA data set dated 21 February 2008. The data set captures relationships between 27184 ASes. We perform two simulation runs. In the first run, YouTube (AS 36561) announces its prefix 208.65.152.0/22. In the subsequent run, PTCL (AS 17557) additionally advertises its sham ownership of the more specific prefix 208.65.153.0/24 to hijack YouTube-addressed traffic. The C-BGP simulations reveal the complete success of the PTCL hijacking attempt: as a result of announcing the longer prefix, PTCL starts receiving all YouTube-addressed traffic with no continued delivery to YouTube itself. This simulation outcome is consistent with the historical accounts of the actual hijacking incident [21]. After each of the runs, C-BGP identifies exactly 100 ASes as being unable to reach any announced prefix. Hence, the number of BGP-connected ASes in the reported simulations stands at 27084.

A. Connectivity of transit ASes to the advertised prefix

In this section, we examine the impact of the hijacking on the BGP connectivity of all transit providers to the announced prefix, i.e., we focus on transit ASes which forward traffic from other ISPs to the advertising entity. Using C-BGP, we determine all converged BGP paths in the simulated scenario. Then, for each transit AS, we count the number of the paths from other ASes through this transit AS to the advertising entity. Below, we interchangeably refer to this value as the number of served BGP paths or *BGP path count* of the transit AS.

Before the prefix hijacking by PTCL, the converged routing involves 2878 transit ASes. For each of the transit ASes, Figure 1 depicts the number of BGP paths served through this AS to YouTube (AS 36561 advertising 208.65.152.0/22) prior to the hijacking. The traditional 16-bit AS numbering space contains 65536 numbers but some ranges of the AS numbers are reserved or yet to be assigned. In Figure 1, these ranges appear with no paths associated with them. One such range covers the numbers from 44587 to 64511, which are not assigned to any AS according to the used CAIDA data set. For some existing ASes, Figure 1 shows no associated paths because these ASes either are unable to reach the advertised prefix 208.65.152.0/22 as per C-BGP or do not serve any transit BGP paths. While the numbers from 64512 to 65534 are designated for private purposes, some of the private AS numbers are represented in Figure 1 with single-path counts. The single-path profile is also common for many regular ASes with a public number. These are the ASes that support a transit path to YouTube for only one of the other 27082 BGP-connected ASes (the overall 27084 ASes minus the served AS minus YouTube). On the opposite side of the connectivity spectrum, AS 11164 – which belongs to TransitRail – serves 9733 paths and is the largest last-hop aggregator of YouTube-bound traffic. Level3 (AS 3356), Hurricane (AS 6939), SprintLink (AS 1239), and Cogent (AS 174) are also connected directly to YouTube and constitute the next four biggest carriers of its incoming traffic with 3983, 3863, 3487, and 3395 served BGP paths respectively.

After PTCL hijacks all YouTube-bound traffic, the number of transit ASes decreases to 2760. For each of the 2760 transit ASes, Figure 2 shows the number of BGP paths served through this AS to PTCL after the hijacking. In comparison

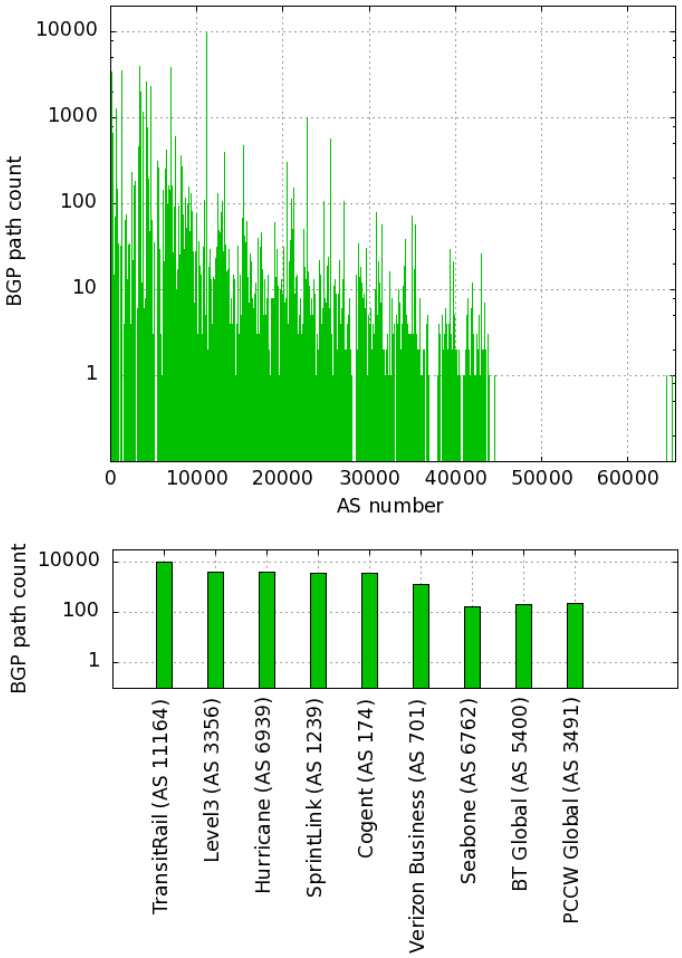

Fig. 1. Number of BGP paths through transit ASes to YouTube (AS 36561 advertising 208.65.152.0/22) before the prefix hijacking by PTCL.

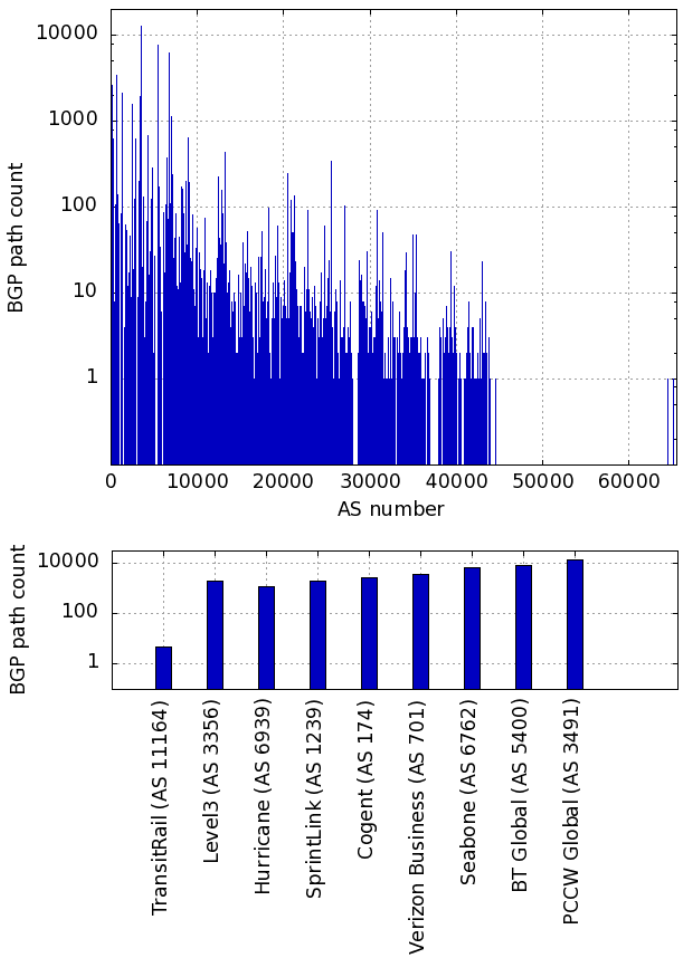

Fig. 2. Number of BGP paths through transit ASes to PTCL (AS 17557 advertising 208.65.153.0/24) after PTCL hijacks all YouTube-bound traffic.

to the routing before the hijacking, the total for the BGP path counts of transit ASes increases by 7.4% from 64213 to 68962. Figure 3 plots the cumulative distributions of the BGP path counts of the transit ASes before and after the hijacking. The plotted curves are quite similar. The share of the transit ASes with exactly one served BGP path is 43.8% before the hijacking and 43.2% after the hijacking. The percentage of the transit ASes with at most 10 served BGP paths is 86.0% before the hijacking and 86.8% after the hijacking. The largest deviation between the routing profiles occurs in the range between 10 and 40 served BGP paths. For example, the 90th percentile of the transit ASes serves at most 16 BGP paths before the hijacking and at most 14 BGP paths after the hijacking. The simulation results confirm the hierarchical structure of the Internet core where a relatively small group of large providers delivers traffic between the other transit ASes. Our results also quantify the following changes in the Internet routing profile. With the 27083 paths to the advertising entity from the other BGP-connected ASes, the hijacking by PTCL extends the average AS-level path length from 3.4 hops to 3.6 hops. The hijacking also increases the average BGP path count (defined as the ratio of the total for the BGP path counts of all transit ASes to the number of such ASes) from 22 to 25.

While the hijacking does not significantly alter the overall distribution of the BGP path counts, some transit providers experience substantial changes in the number of served BGP paths. For example, the BGP path counts for the five giants TransitRail, Level3, Hurricane, SprintLink, and Cogent shrink from 9733, 3983, 3863, 3487, and 3395 to 5, 1915, 1124, 1854, and 2648 respectively. The preserved paths do not lead to YouTube anymore but instead contribute to the hijacking success of PTCL. The providers with the largest BGP path counts after the hijacking are PCCW Global (AS 3491), BT Global (AS 5400), and Telecom Italia Seabone (AS 6762): the hijacking boosts their BGP path counts from 219, 206, and 164 to 12942, 7793, and 6319 respectively. Figure 4 summarizes the changes in the BGP path counts of transit ASes as a consequence of the hijacking. 895 ASes decrease their BGP path counts with the cumulative loss of 32430 paths, and only 481 ASes increase their BGP path counts with the aggregate gain of 37179 paths. Figure 5 highlights the 5 largest losers and 5 biggest winners of transit paths as a result of the

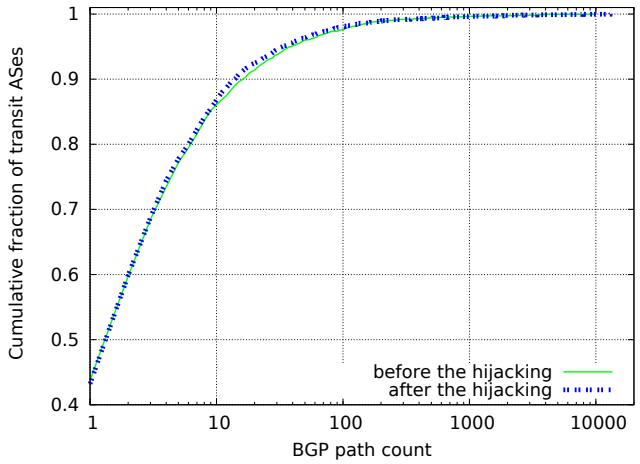

Fig. 3. Cumulative distribution of the BGP path counts of transit ASes before and after the prefix hijacking by PTCL.

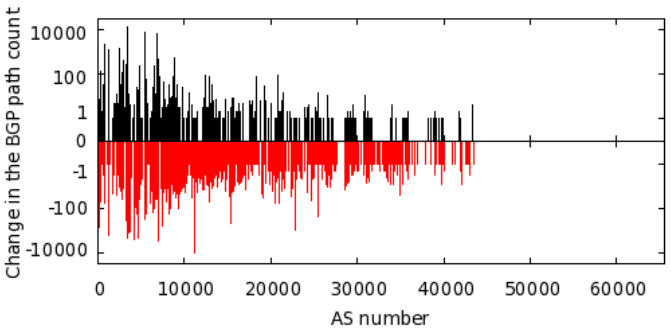

Fig. 4. Change in the BGP path count as a consequence of the prefix hijacking by PTCL.

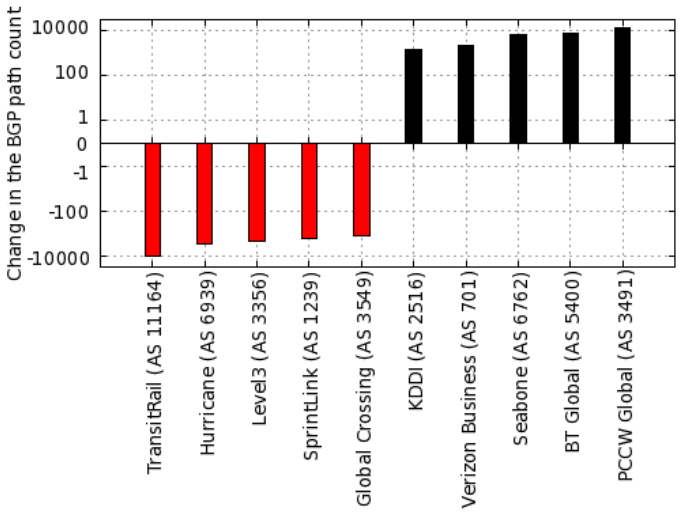

Fig. 5. Change in the BGP path count for the 5 largest losers and 5 biggest winners of transit paths as a consequence of the prefix hijacking.

prefix hijacking: TransitRail, Hurricane, Level3, SprintLink, and Global Crossing (AS 3549) lose respectively 9728, 2739, 2068, 1623, and 1163 transit paths while KDDI (AS 2516), Verizon Business (AS 701), Telecom Italia Seabone, BT Global, and PCCW Global gain respectively 1336, 2152, 6155, 7587, and 12723 transit paths.

B. Traffic on inter-provider links

Whereas the payments of providers depend on the volumes of traffic communicated over the corresponding inter-provider links, this section utilizes the computed BGP paths to derive the link loads of all inter-AS links. In either of the simulations, before or after the prefix hijacking by PTCL, we consider only YouTube-addressed traffic and transmit a portion of such traffic from each of the 27084 BGP-connected ASes. Although the goal of the actual prefix-hijacking incident was to blackhole YouTube-addressed traffic, the hijack-to-discard version of prefix hijacking is not sustainable in the long run, and we are mostly interested in hypothetical long-term hijack-to-deliver instances where the hijacker forwards the intercepted traffic to the destination. Under symmetric routing and proportionality of incoming and outgoing traffic volumes, one could use our unidirectional traffic methodology to determine the inter-AS link loads (and thereby the ISP payments) for bidirectional hijack-to-deliver scenarios. It would be even more desirable to simulate the bidirectional scenarios directly. However, the scalability properties of C-BGP do not make it feasible to complete simulations where each of the 27084 ASes announces a separate prefix. For the same C-BGP scalability reasons, we do not simulate the Internet cross-traffic, i.e., the traffic that both starts and terminates in other ASes than YouTube and PTCL.

While YouTube-bound traffic contains video clips uploaded by YouTube users as well as requests for clip downloads, the uploads are likely to dominate the requests in terms of the traffic volume, and we focus only on this former type of traffic. According to [32], video was uploaded to YouTube in 2008 at the rate of 12 hours per minute, meaning that the volume of video clips uploaded every minute was such that playing them one after another would take 12 hours. After analyzing a collection of video clips in the FLV (FLash Video) format with playing times in the range from 1 minute to 1.3 hours, we estimate that 1 hour of playing time corresponds to 100 MB of data. Hence, our estimate for the average year-2008 rate of YouTube-addressed traffic is 160 Mbps.

Determining the origins of the YouTube-addressed traffic is a more challenging task. ISPs commonly perceive exchanged traffic volumes as sensitive information. While we are aware of anonymized data sets that quantify the relative potency of various ASes to generate traffic, the goal of our study necessitates associating a generated traffic volume with each specific AS. Without having access to real data sets of the latter type, we allocate the generated traffic to all BGP-connected ASes uniformly, i.e., each AS in our C-BGP simulations generates YouTube-addressed traffic at the same rate of 6 Kbps.

With the knowledge of the rates at which the ASes inject their traffic toward the announced prefix along the computed

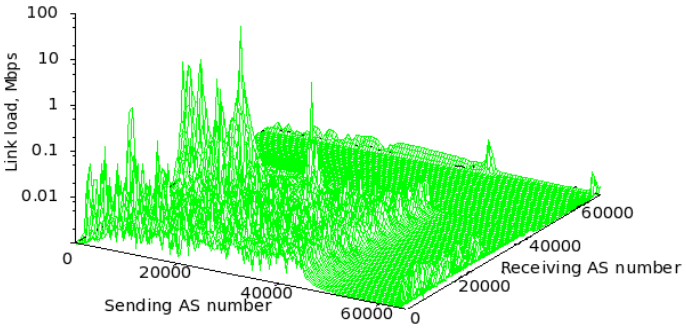

Fig. 6. Link loads of the inter-AS links before the prefix hijacking by PTCL.

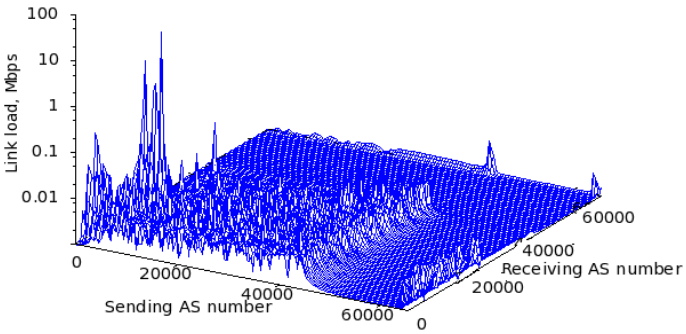

Fig. 7. Link loads of the inter-AS links after the prefix hijacking by PTCL.

BGP paths, we add up the overlapping traffic flows to determine the link load of every inter-AS link. Figure 6 depicts the loads of all inter-provider links before the prefix hijacking by PTCL. The three-dimensional graph handles each inter-AS link as a two-tuple that contains the AS numbers of its sending and receiving end points. To plot the loads for the 65536-by-65536 link space as a 128-by-128 mesh, we use command dgrid3d with norm 3 in gnuplot [33]. Despite the aggregation of individual link loads by gnuplot, Figure 6 reveals interesting features of the overall traffic pattern, e.g., ridges that run in parallel with the Sending-AS-number axis. The peaks on the tallest ridge correspond to the links leading to YouTube (AS 36561) from its last-hop aggregators such as TransitRail (AS 11164, the highest peak), Level3, Hurricane, SprintLink, and Cogent. The second tallest ridge corresponds to the links that lead to TransitRail from its suppliers of YouTube-bound traffic.

Figure 7 shows the changed traffic pattern after the successful hijacking attempt. The tallest ridge runs now toward PTCL (AS 17557) and is dominated by the peaks corresponding to the links that lead to the hijacker from its providers PCCW Global, BT Global, and Telecom Italia Seabone. The hijacking affects not only the links chosen for YouTube-addressed traffic but also the type of the chosen links. For instance, before PTCL hijacks the delivery, Turk Telekom (AS 9121) routes through its peer Swisscom (AS 3303), the Turk Telekom traffic constitutes around 34% of the total traffic received by Swisscom, and Swisscom passes the traffic to its customer FLAG Telecom (AS 15412). After the hijacking, Swisscom routes through its peer PCCW Global without opening this path for Turk Telekom, and Turk Telekom ends up routing through its provider Telecom Italia Seabone.

C. Payments of providers

At the last step of our three-step evaluation method, we now calculate payments of providers. To achieve this, we utilize Equations 1, 2, and 3, link types from the CAIDA data set, and link loads from Section III-B.

Figure 8 shows the payments of all the 27084 BGP-connected ASes before PTCL hijacks the YouTube-addressed traffic. The simulations show negative payments for 25796 ASes, i.e., 95% of all the BGP-connected AS population. This observation is again consistent with the hierarchical structure of the Internet: most ASes pay money to their transit providers but provide no transit services for any AS customer and thus need to recover their transit costs by charging their individual users. YouTube (AS 36561) has the largest negative payment of $823, comprising about $2 spent on peering with RETN (AS 25462), Net Access (AS 8001), and Rogers Cable (AS 812) as well as $254, $130, $127, $117, $115, and $78 paid respectively to its six providers TransitRail, Level3, Hurricane, SprintLink, Cogent, and Global Crossing. YouTube recovers the above traffic costs indirectly, by exploiting the video that the traffic delivers. More interestingly, TransitRail (AS 11164), REACH Global (AS 4637), China Telecom (AS 4134), and Limelight Networks (AS 22822) also experience negative payments of $175, $26, $20, and $2 respectively. This outcome is surprising because these four transit ASes are expected to make money, rather than to lose money, on the transit services that they provide. The outcome is connected to ignoring the Internet cross-traffic. For instance, TransitRail receives YouTube-addressed traffic from its 13 providers – such as REACH Global, China Telecom, Limelight Networks, and LG Dacom (AS 3786) – and passes the aggregated traffic to its customer YouTube. Due to the sublinear pricing, TransitRail gets from YouTube $176 less for the aggregate traffic than the total of $430 paid by TransitRail to its 13 providers for their individual contributions (TransitRail also recovers $1 from its customers that route their YouTube-addressed traffic through TransitRail). With the Internet cross-traffic included, the link from TransitRail to YouTube would carry less traffic than the link from either of the 13 providers to TransitRail, and the per-unit price of YouTube-addressed traffic for these links would be smaller than for the link to YouTube, not larger as in the reported evaluation. In the future, we will address this limitation and model the cross-traffic as well. Out of the 25796 ASes in the red, 23773 ASes pay exactly $0.26 each, which corresponds to 6 Kbps of traffic injected by a non-transit AS. There are also 1288 ASes in the black. The largest positive payments of $720, $706, $637, $357, and $289 are respectively achieved by Level3 (AS 3356), SprintLink (AS 1239), Cogent (AS 174), Verizon Business (AS 701), and Global Crossing (AS 3549). The sum of all positive payments is $7109, the

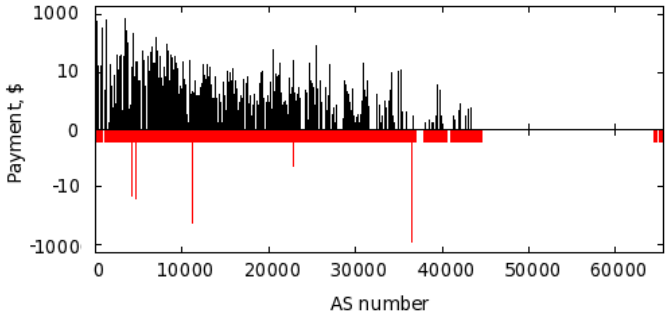

Fig. 8. Payments of all ASes before PTCL hijacks the traffic addressed to YouTube (AS 36561).

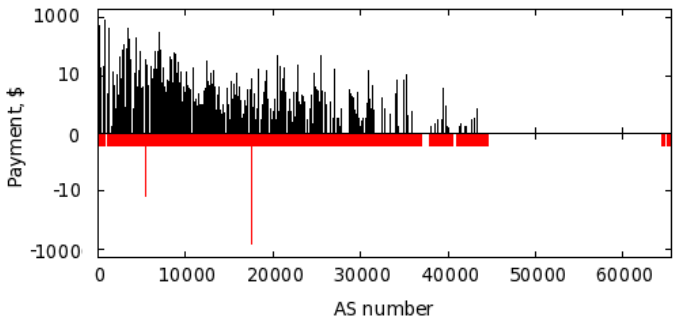

Fig. 9. Payments of all ASes after the hijacking of the YouTube-addressed traffic by PTCL (AS 17557).

total of the negative payments is $7491, the difference of $382 represents the peering expenses of the ASes.

Figure 9 depicts the AS payment distribution after the prefix hijacking. 25889 ASes experience negative payments with the total of $7203, 1194 ASes attract positive payments with the total of $6872, and $331 is spent on peering. PTCL (AS 17557) has the largest negative payment of $706: while attracting around $6 from its 24 customers, PTCL racks up the highest bill by spending $314, $215, and $183 respectively on the traffic received from its providers PCCW Global, BT Global, and Telecom Italia Seabone. The only other AS with a considerable negative payment is BT Global (AS 5400), which ends up with $14 in the red. We attribute this negative payment to the same reason as for TransitRail before the hijacking, i.e., to ignoring the cross-traffic. Verizon Business (AS 701), Cogent (AS 174), SprintLink (AS 1239), Level3 (AS 3356), and AT&T WorldNet (AS 7018) are the providers with the five biggest positive payments of $835, $571, $437, $423, and $300 respectively.

Figure 10 displays the changes in the payments of all ASes as a consequence of the prefix hijacking. 1036 ASes decrease their payments with the cumulative loss of $2623, and only 452 ASes increase their payments with the aggregate gain of $2674. Figure 11 focuses on PTCL, YouTube, and the other 5 largest financial losers and 5 biggest financial winners: PTCL, Level3, SprintLink, Global Crossing, BT Global, and LG Dacom lose $708, $297, $269, $109, $70, and $69 respectively while Verizon EMEA (AS 702), TeliaSonera (AS 1299), AT&T WorldNet, TransitRail, Verizon Business, and YouTube gain $87, $130, $137, $176, $476, and $823 respectively. Whereas the non-transit ASes spend more than $6000 on the transit of their YouTube-addressed traffic, the transit ASes are financially interested in the choice of inter-AS links for the traffic. Comparing Figure 11 with Figure 5, we see that the types of inter-AS links (and not accounting for the cross-traffic) are very influential in translating the BGP path counts into the financial fortunes of the giant transit providers. Although Level3, SprintLink, and Global Crossing are among the largest losers – and Verizon Business remains a big winner – with respect to both metrics, BT Global increases its BGP path count but loses money, and TransitRail financially benefits from serving less paths.

As Figures 10 and 11 demonstrate, PTCL makes itself the largest financial loser: the hijacking reverts its positive payment of $2 to the negative payment of $706. Do the benefits of PTCL from diverting to itself the YouTube-addressed traffic justify the associated price tag? Also, while YouTube saves the record $823 by not receiving its traffic, the hijacking results in the loss of the incoming traffic and thereby prevents YouTube from earning an even larger income on the video that the traffic carries. Do YouTube and the transit providers that lose transit traffic have effective means to shut off or at least alleviate the prefix hijacking? Below, we explore these and other related questions in the context of a broader discussion on viability of hijacking-based traffic attraction.

IV. Viability of Traffic Attraction via Hijacking

Viability of hijacking-based traffic attraction is a multi-dimensional topic. Without pretending to be comprehensive, this section examines the implications of the above results along some of the dimensions, including incentives, technical feasibility, legal and business considerations.

Does any provider derive sufficiently high financial benefits from the hijacking-based traffic attraction? Our results for the hijacking of the YouTube-addressed traffic show that the biggest winners gain on the order of $100 per month. Such a reward does not seem a strong incentive for the transit giants. However, Section III studies the hijacking of only one prefix, and the payment change for Verizon Business – the biggest winner in this regard among the transit ASes – is high in relative terms, 58% of the total paid by YouTube before the hijacking. If a large number of popular prefixes is hijacked, the cumulative reward can be significant and impel the benefited ASes to support the hijacking.

What is the incentive for the initiator of the prefix hijacking? In Section III, PTCL is $2 in the black before the hijacking (thanks to the payments from own customers for the transit of their YouTube-bound traffic) but finds itself $706 in the red after the hijacking. While the objective of PTCL in the actual hijacking incident was to block access to YouTube, long-term attraction and discard of unwanted traffic do not appear sustainable in the long run. Another intriguing option

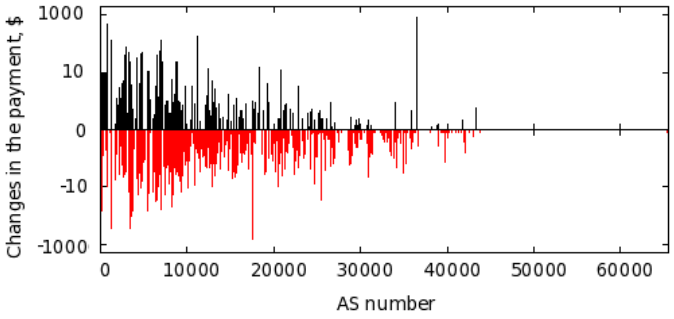

Fig. 10. Changes in the payments of all ASes as a consequence of the prefix hijacking by PTCL.

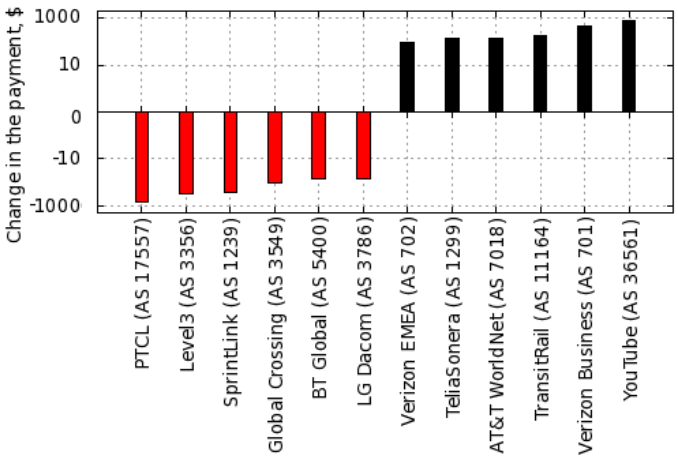

Fig. 11. Changes in the payments of PTCL, YouTube, and the other 5 largest financial losers and 5 biggest financial winners as a consequence of the prefix hijacking.

for PTCL is to imitate YouTube, i.e., to recover the traffic costs by exploiting the video that the hijacked traffic delivers. While fake Nikes are a big business, we do not investigate whether a fake YouTube exists or can financially sustain itself in principle. Instead, we focus on transit payments, the main source of revenue considered in our paper. PTCL is an obvious loser in this regard. We attribute the negative payment to the role of Pakistan Telecom in the Internet transit hierarchy: PTCL is a small player on the global transit scale, and the flow of the hijacked traffic is such PTCL ends up paying much more to own providers than it collects from its customers. As Section III confirms, the biggest financial winners among transit ASes are large ISPs, e.g., top-tier provider-free ISPs that exchange traffic with their neighbors without making transit payments to any of them. Because of the strongest financial incentives, large transit ISPs are the most likely initiators of hijacking-based traffic attraction.

Do the prefix owners or transit ASes that lose revenue due to hijacking have effective means to shut off or at least alleviate the prefix hijacking? Even detection of the hijacking is not a straightforward task when the intercepted traffic is subsequently delivered by the hijacker to the prefix owner. There exist tools for monitoring BGP announcements and IP forwarding paths but their effectiveness is limited. Besides, the prefix owner does not have strong incentives to be concerned about a hijack-to-deliver detour of its traffic unless the traffic diversion significantly disrupts the quality of the path or increases the traffic expenses of the prefix owner.

Stopping the hijacking-based traffic attraction altogether is quite difficult, especially for a non-transit owner of the hijacked prefix. As long as all ASes on the path from a traffic origin through the hijacker to the prefix owner are comfortable with attracting the traffic to this new path, the prefix owner or ASes off the new path do not have assured means for reverting the delivery to the path expected under the traditional BGP announcement practices. The affected parties can attempt and succeed in hijacking the hijacked delivery back, e.g., by announcing a more specific prefix as YouTube did in the real prefix-hijacking incident. Thus, hijacking-based traffic attraction has a potential for creating a fertile ground for tussles between providers [34]. The announcements of more specific prefixes fragment the prefix spaces, increase the forwarding tables, and thereby jeopardize the scalability of Internet routing. On the other hand, some researchers believe that Internet-scale routing with flat addresses is technically feasible [35], [36]. It remains to be seen whether hijacking-based traffic attraction will lead to massive inter-provider tussles and pose a danger to routing scalability.

Beyond the technical countermeasures, the afflicted parties can try to neutralize hijacking-based traffic attraction through litigation. The global nature of the Internet complicates the judicial process for prefix-hijacking cases. In particular, the hijacker and prefix owner – as exemplified by Pakistan Telecom and YouTube in the real prefix-hijacking incident – can operate under different national legal systems. While the body of laws governing the Internet is generally slim but growing, we are not aware of any precedents of prefix-hijacking litigation. Prior laws and guidelines issued by various governments for other Internet-related disputes demonstrate that the outcomes of the legal battles are highly unpredictable. More specifically, it is hard to predict whether the contested actions deviating from an expected Internet behavior will be ruled illegal or legitimate, or even deserving a special protection by the law. Some national governments take steps – e.g., in the context of network neutrality [37], [38] – to curb the technical means that transit ISPs have for managing the served traffic. On the other hand, peer-to-peer systems [39] and content distribution networks [40], which represent dramatic departures from traditional ways to disseminate information over the Internet, have enjoyed a wide adoption around the world despite disputes about lawfulness of such distribution methods.

With limited options in the technical and legal spheres, the business side of the Internet is likely to serve as an important arena for settling the tussles of hijacking-based traffic attraction. In order to offer the universal Internet connectivity to own customers, any ISP anywhere in the routing hierarchy

has to maintain business relationships with other ASes. While different ASes have clearly different negotiation power, losing a customer or peer is rarely a desirable outcome for the prefix-hijacking ISP. In the business world, reputations are tangible assets: a bad reputation can severely diminish the ability of the ISP to negotiate transit and peering contracts. Hence, if the ISP community as a whole starts to deem hijacking-based traffic attraction unacceptable, the risk of a bad reputation can serve as a strong disincentive for an ISP to boost revenues through hijacking-based traffic attraction.

V. Future research directions

To the best of our knowledge, this paper is the first to present an economic perspective on prefix hijacking as a means for revenue boosting via traffic attraction. While the paper – by using the real data to drive the simulations – offers interesting insights, it also identifies promising directions for further research of the topic. First, our results suggest that large transit ISPs, such as top-tier provider-free ASes, have the strongest financial incentives to hijack a prefix for hijack-to-deliver traffic attraction. Hence, a future study should directly examine scenarios where a giant transit ISP hijacks popular prefixes to intercept the corresponding traffic and then deliver the intercepted traffic to the proper destinations. Second, our analysis reveals that accounting for the Internet cross-traffic is highly important for accurate translation of the BGP path counts into the payments of providers. Therefore, future studies should strive to model the traffic matrices realistically, not only for the hijacked prefix but also for the cross-traffic. Third, we observe that the default valley-free configuration of C-BGP can yield paths that are different from those reported for the actual hijacking of the YouTube prefix by PTCL. Consequently, the simulator needs to be enhanced with BGP import and export policies of a better fidelity. Fourth, this paper indicates that hijacking-based traffic attraction can create a fertile ground for tussles between ISPs, in both technical and business spheres. Thus, hijacking-driven inter-provider tussles represent a promising area for future work. As our experience shows, availability of real data and scalability of C-BGP simulations are serious obstacles for the above lines of future research. To overcome these challenges, one might need to enhance our data-driven simulation methodology with additional modeling and analysis.

VI. Related work

This section briefly discusses prior work related to prefix hijacking and traffic attraction. Unlike our paper which focuses on the economic implications, the prior studies are driven by security considerations and aim at undermining the effectiveness of prefix hijacking. Lad, Massey, et al. [7] develop PHAS (Prefix Hijack Alert System), an online system that notifies the prefix owner when the BGP path to the prefix changes. McArthur and Guirguis [6] explore stealthy forms of prefix hijacking that attract small amounts of traffic and thereby avoid detection. Zhang, Zhao, and Wu [13] investigate an attack where an AS selectively drops BGP announcements to severely disrupt the routing. Goldberg, Halevi, et al. [19] suggest that it is difficult to achieve honesty in BGP announcement practices without reliance on heavyweight forwarding-aware protocols that verify and enforce AS-level paths. Ballani, Francis, and Zhang [4] take a stochastic approach to hijack-to-deliver traffic attraction and quantify the probability of the hijacking success for various types of ASes. The related work agrees with our observations that technical countermeasures against prefix hijacking are limited and that hijacking-based traffic attraction opens plentiful opportunities for inter-provider tussles.

VII. Conclusion

In this paper, we investigated how prefix hijacking impacted payments for inter-AS traffic and whether an ISP could boost its revenue by means of prefix hijacking. To examine prefix hijacking from this novel economic angle, we used real data as a basis for C-BGP simulations. The specific starting point for our study was the actual incident where Pakistan Telecom hijacked the prefix owned by YouTube. We conducted the simulations in the CAIDA topology with 27084 BGP-connected ASes and synthetic traffic demands for the YouTube prefix. Then, we calculated the link loads of all inter-AS links and determined the payments of all providers. Our analysis of the results from technical, business, and legal perspectives suggested that hijacking-based traffic attraction was a viable strategy that could create a fertile ground for tussles between ISPs. In particular, top-tier provider-free ISPs appeared to have the strongest financial incentives for engaging in hijack-to-deliver attraction of traffic addressed to popular prefixes. We also discussed directions for future research in the area of hijacking-based traffic attraction.

References

[1] Y. Rekhter and T. Li, "A Border Gateway Protocol (BGP-4)," *RFC 1771*, March 1995.

[2] Information Sciences Institute, University of Southern California, "Internet Protocol DARPA Internet Program Protocol Specification," *RFC 791*, September 1981.

[3] S. Dharmapurikar, P. Krishnamurthy, and D. Taylor, "Longest Prefix Matching Using Bloom Filters," *IEEE/ACM Transactions on Networking*, vol. 14, no. 2, pp. 397–409, April 2006.

[4] H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," *In Proceedings of ACM SIGCOMM 2007*, pp. 265–276, August 2007.

[5] L. Li and C. Chen, "Exploring Possible Strategies for Competitions between Autonomous Systems," *In Proceedings of IEEE ICC 2008*, pp. 5919–5923, May 2008.

[6] C. McArthur and M. Guirguis, "Stealthy IP Prefix Hijacking: Don't Bite Off More Than You Can Chew," *In Proceedings of IEEE GLOBECOM 2009*, pp. 2480–2485, November 2009.

[7] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A Prefix Hijack Alert System," *In Proceedings of USENIX-SS 2006*, pp. 153–166, July 2006.

[8] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis, "A Light-Weight Distributed Scheme for Detecting IP Prefix Hijacks in Real-Time," *In Proceedings of ACM SIGCOMM 2007*, pp. 277–288, August 2007.

[9] X. Hu and Z. M. Mao, "Accurate Real-time Identification of IP Hijacking," *In Proceedings of IEEE SSP 2007*, pp. 3–17, May 2007.

[10] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush, "iSPY: Detecting IP Prefix Hijacking on My Own," *In Proceedings of ACM SIGCOMM 2008*, pp. 327–338, August 2008.

[11] M. Lad, R. Oliveira, B. Zhang, and L. Zhang, "Understanding Resiliency of Internet Topology Against Prefix Hijack Attacks," *In Proceedings of IEEE/IFIP DSN 2007*, pp. 368–377, June 2007.

[12] Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao, "Practical Defenses Against BGP Prefix Hijacking," *In Proceedings of ACM CoNEXT 2007*, pp. 1–12, December 2007.

[13] K. Zhang, X. Zhao, and S. Wu, "An Analysis on Selective Dropping Attack in BGP," *In Proceedings of IEEE IPCCC 2004*, pp. 593–599, April 2004.

[14] K. Butler, T. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, January 2010.

[15] O. Nordstrom and C. Dovrolis, "Beware of BGP Attacks," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 2, pp. 1–8, April 2004.

[16] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 4, pp. 582–592, April 2000.

[17] F. Sanchez and Z. Duan, "Region-based BGP announcement filtering for improved BGP security," *In Proceedings of ACM ASIACCS 2010*, pp. 89–100, April 2010.

[18] "RIPE Routing Registry." [Online]. Available: http://www.ripe.net/db/irr.html

[19] S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R. N. Wright, "Rationality and Traffic Attraction: Incentives for Honest Path Announcements in BGP," *In Proceedings of ACM SIGCOMM 2008*, pp. 267–278, August 2008.

[20] R. Teixeira, A. Shaikh, T. Griffin, and J. Rexford, "Dynamics of Hot-Potato Routing in IP Networks," *SIGMETRICS Perform. Eval. Rev.*, vol. 32, no. 1, pp. 307–319, June 2004.

[21] "YouTube Hijacking: A RIPE NCC RIS case study," February 2008. [Online]. Available: http://www.ripe.net/news/study-youtube-hijacking.html

[22] "NANOG Mailing List," 2008. [Online]. Available: http://www.nanog.org/mailinglist/

[23] "RIPE RIS." [Online]. Available: http://www.ripe.net/ris/

[24] "BGPmon." [Online]. Available: http://bgpmon.net/

[25] B. Quoitin and S. Uhlig, "Modeling the Routing of an Autonomous System with C-BGP," *IEEE Network Magazine*, vol. 19, no. 6, pp. 12–19, November 2005.

[26] S. Qiu, P. McDaniel, and F. Monrose, "Toward Valley-Free Inter-domain Routing," *In Proceedings of IEEE ICC 2007*, pp. 2009–2016, June 2007.

[27] B. Quoitin, "CBGP, Sourceforge, SCM Repositories," July 2010. [Online]. Available: http://cbgp.svn.sourceforge.net/viewvc/cbgp?revision=1131\&view=revision

[28] "CAIDA." [Online]. Available: http://www.caida.org/data/active/as-relationships/index.xml

[29] A. Dhamdhere and C. Dovrolis, "Can ISPs be Profitable Without Violating Network Neutrality?" *In Proceedings of NetEcon 2008*, pp. 13–18, August 2008.

[30] H. Chang, S. Jamin, and W. Willinger, "To Peer or not to Peer: Modeling the Evolution of the Internets AS-level Topology," *In Proceedings of IEEE INFOCOM 2006*, pp. 1–12, April 2006.

[31] "DrPeering International," August 2010. [Online]. Available: http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php

[32] "Oops, pow, surprise...24 hours of video all up in your eyes!" March 2010. [Online]. Available: http://youtube-global.blogspot.com/2010/03/oops-pow-surprise24-hours-of-video-all.html

[33] "Gnuplot." [Online]. Available: http://www.gnuplot.info/

[34] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," *In Proceedings of ACM SIGCOMM 2002*, pp. 347–356, August 2002.

[35] M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, and I. Stoica, "ROFL: Routing on Flat Labels," *In Proceedings of ACM SIGCOMM 2006*, pp. 363–374, September 2006.

[36] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A Data-Oriented (and Beyond) Network Architecture," *In Proceedings of ACM SIGCOMM 2007*, pp. 181–192, August 2007.

[37] J. Crowcroft, "Net Neutrality: The Technical Side of the Debate ˜ A White Paper," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 1, pp. 49–56, January 2007.

[38] M. Yuksel, K. K. Ramakrishnan, S. Kalyanaraman, J. D. Houle, and R. Sadhvani, "Class-of-Service in IP Backbones: Informing the Network Neutrality Debate," *In Proceedings of ACM SIGMETRICS 2008*, pp. 435–436, June 2008.

[39] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications," *In Proceedings of ACM SIGCOMM 2001*, pp. 149–160, August 2001.

[40] B. Krishnamurthy, C. Wills, and Y. Zhang, "On the Use and Performance of Content Distribution Networks," *In Proceedings of ACM SIGCOMM Internet Measurement Workshop 2001*, pp. 169–182, November 2001.