Exploiting Overlays for Practical and Profitable Traffic Attraction

Jeff Seibert
Purdue University
jcseiber@cs.purdue.edu

Abstract— **The Internet is heading towards efficiency with the continued deployment of technologies such as Content Distribution Networks (CDN) and with the introduction of Peer-to-Peer (P2P) localization. These technologies reduce the amount of costly traffic that Internet Service Providers (ISPs) must send to other ISPs. While these advances will undoubtedly cause many residential ISPs to gain profit, the transit ISPs that carry traffic for them will actually lose revenue.**

In this work, we identify different means by which transit ISPs can increase revenue by subverting CDN and P2P systems and thus causing traffic to flow through them on profitable paths. We evaluate how profitable such techniques are by collecting datasets from the popular BitTorrent P2P system and the worldwide Akamai and YouTube CDNs. We find that many ISPs are able to at least double their revenue-generating P2P traffic and some are also able to attract significant amounts of CDN traffic.

I. Introduction

Transit ISPs provide the service of carrying traffic between residential ISPs and content providers on the Internet. Their business model also captures this such that, generally, the more traffic they carry, the larger their profits will be. Residential ISPs, who provide Internet service for end-users, and also content providers, who generate content for consumption by end-users, therefore seek to reduce their costs of operation by minimizing the amount of traffic that leaves their networks.

Content providers often reduce the costs associated with carrying traffic by using CDNs. A CDN is an overlay network that has replica servers placed in many locations that are typically close to the end-user. Then whenever content is requested by someone, he will be redirected to a replica that has close locality to it, which is usually determined by latency. Thus the more replicas deployed and the more similar content is requested, the less transit ISPs are used and the less profit they make.

Residential ISPs, on the other hand, have typically seen P2P systems as a nuisance for economic reasons. While P2P systems often provide similar functionality as a CDN, providing Internet-scale distribution of content, they often use an exorbitant amount of bandwidth and can unnecessarily increase the amount of traffic sent through a transit ISP. This is due to P2P systems being oblivious to other nodes that are also inside the same ISP, thus multiple copies of the same content are downloaded from outside the ISP. Recently though, many works have proposed P2P localization schemes to make P2P systems more "ISP friendly" [1]–[5]. These localization schemes allow nodes in the same ISP to find and trade data with each other, rather than having to receive multiple copies from other ISPs. These techniques have the potential to transform P2P systems into effective content distribution systems in terms of performance and costs to the ISP.

In this work, we investigate ways in which the mechanisms used by CDNs and P2P systems to actually foster locality can be exploited to increase traffic flowing on profitable paths in a transit ISP. Then, through measurements of the Akamai and YouTube CDNs, as well as the BitTorrent P2P system, we analyze how effective traffic attraction can be in practice.

The structure of this paper is as follows: We give an overview of proposed mechanisms for providing locality in CDNs and P2P systems in Section II and then identify how to use them for traffic attraction in Section III. We describe our datasets and present the results of our analysis in Section IV. Finally, we describe future work and conclude in Section V.

II. Background

CDNs and P2P systems have different mechanisms that they employ to help users find replicas of content that are close to them. CDNs typically use DNS redirections to direct users to content. Specifically, when a user wants to download content, the URL might refer to a domain for which the corresponding IP address is unknown. To find the IP address, the user queries its recursive DNS (RDNS) server with the domain name. If the RDNS server does not know, it will do a recursive DNS query, asking the CDN's authoritative nameservers. This nameserver will respond with an IP address of a machine that has the content cached on it. Since the CDN has machines with the same content in multiple places, it chooses one that is close to the RDNS server. The RDNS server then returns the results to the user who can finally download the content from the CDN replica.

P2P systems, on the other hand, have had a number of different proposed ways to foster locality. We now describe these different means:

P2P-only solutions: P2P systems can unilaterally attempt to localize the traffic that they generate without any intervention from ISPs. Ono [3] is one such example where nodes try to discern if another node is close by and thus in the same ISP by if they have similar CDN (e.g. Akamai) redirections. P2P systems, such as BitTorrent [6], could also modify their trackers, so that instead of a random selection of nodes, a localized selection is given.

ISP and P2P cooperation: Some research has suggested that P2P systems can work together with ISPs for mutual benefit. This is due to P2P systems generally performing better when biasing towards local nodes and residential ISPs decrease costs as a result. Some have suggested protocols by which P2P systems can query oracles set up by ISPs to find out who the ISP would prefer the P2P user to download content from [1], [2]. In these scenarios the ISP has direct influence over who the P2P user downloads from.

ISP-only solutions: ISPs can also employ techniques to localize traffic in its own networks. Bindal et. al. [5] suggest that ISPs employ deep packet inspection to keep track of BitTorrent peers inside of the ISP and to rewrite packets so that peers biasedly learn about those inside the ISP. They also suggest that ISPs set up HTTP proxies that intercept the HTTP traffic that occurs between BitTorrent trackers and peers.

III. Traffic Attraction

We first describe different scenarios in how traffic attraction can be used and second how to use the techniques described in Section II for traffic attraction.

A. Types of Attraction

We label ISPs as being residential, transit, or both. Transit ISPs carry traffic for other ISPs, hence they are not necessarily the originators or destination of traffic. Residential ISPs have actual users, thus they are the originators or destination of traffic. We note that an ISP's type will allow it to do different types of traffic attraction. We also note that while a strictly residential ISP can cause traffic attraction, they will not directly benefit from it, as they always have to pay other ISPs to send or receive traffic. We now delineate possible cases of traffic attraction. We note that both the CDN and P2P mechanisms that can cause traffic attraction, as will be described in Section III-B and III-C, can cause any of these particular cases. For ease of exposition we use Figure 1 and assume that CDN replica nodes are only found in residential ISPs, although in practice they can be found in other types.

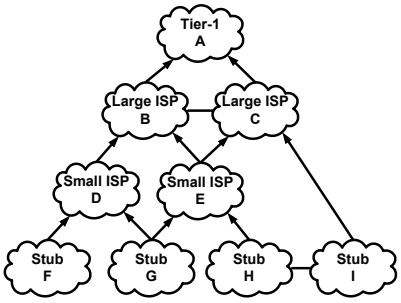

Fig. 1. A hierarchical view of the Internet. Transit ISPs provide service to smaller ISPs through revenue-generating customer-provider links (denoted with a directed edge) or carry traffic for free through a peering link (denoted with an undirected edge).

Transit Attraction: ISP B is a transit ISP and ISPs D and E are residential ISPs. ISP B causes nodes in ISP E to exchange traffic with nodes in ISP D. Thus traffic will flow through ISP B. Any transit ISP can do this attraction.

Residential Attraction: ISP B is both a transit and residential ISP. ISP D is a residential ISP. ISP B causes its own nodes to exchange traffic with those in ISP D. Only ISPs that are both residential and transits can do this attraction.

Congestion Attraction: ISPs D and E are residential ISPs and ISP C is any kind of ISP. ISP C causes nodes in ISP D to exchange traffic with nodes in ISP E. ISP C does not provide transit service, as there is either a direct link or another ISP provides transit service. An ISP may launch this type of attack if it wishes to create congestion in a network (e.g., ISP B) that it does not own. Any kind of ISP can launch this attack.

B. CDN Traffic Attraction

DNS cache poisoning: In this well-known attack the adversary tricks an RDNS into caching an IP address of its own choosing for a particular domain name. To trigger the poisoning, the attacker sends to the RDNS a query for a domain name, which if it does not exist in the cache, the RDNS sends a query to an authoritative nameserver. The attacker must then send a well-crafted reply to the RDNS, before the authoritative nameserver actually responds. To accomplish this, the attacker must try to guess the 16-bit transaction ID that is included in the query, and depending on the implementation, also the UDP source port used by the RDNS. Recent research has also suggested using wildcard domain names to increase the security of DNS by making adversaries guess an even longer string [7]. However, these schemes do not attempt to protect against an adversary that can actually see the DNS traffic. ISPs then are in prime positions to take advantage of DNS cache poisoning.

Rewrite DNS responses: Alternatively, if a DNS response for a CDN is detected going over a transit ISP, the ISP could rewrite the response so that it contains the IP addresses of the CDN replica it desires.

Note that for the previous examples DNSSEC will not protect against these attacks. The CDN owner will digitally sign each of the IP addresses that it owns. However, as the adversary is simply exchanging one machine owned by the CDN for another, he can collect the signed portion of the DNS responses for the different replica's IP addresses and then replay that to the RDNS server of its choosing.

Deploy RDNS: If the ISP has its own RDNS, it can return CDN replicas that will cause its own users to draw traffic from them on profitable paths. Furthermore, if the RDNS is open to queries from anyone on the Internet, it can choose replicas based on the querying IP address.

Rewrite webpages: If the webpage is transmitted over a transit ISP, it could rewrite the domain name of the content to the IP address of the CDN replica that it desires.

C. P2P Traffic Attraction

Deploy Tracker: As anyone can deploy a tracker such as those used in BitTorrent, and P2P protocols are usually well-known, an ISP can deploy its own modified version. Trackers usually hand out a random selection of nodes to whoever asks for some. However, with its own tracker, an ISP can know

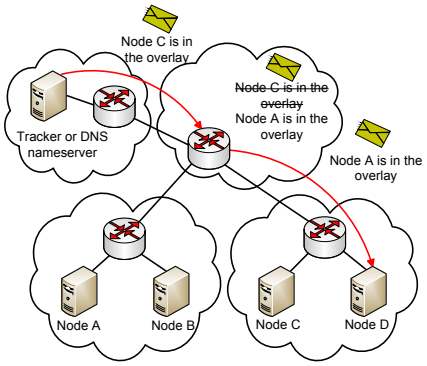

Fig. 2. An example of how traffic attraction could take place. Each cloud is a different ISP, where one ISP provides transit to the rest. The transit ISP changes the packet sent by the tracker (or DNS server) so that resulting traffic will flow through the transit ISP.

all the participants who want some particular content and thus can easily direct neighbor selection to its own preferences.

Oracles: As stated earlier, recent research has introduced the notion of using oracles as ways for ISPs to facilitate localization [1], [2]. Both trackers and nodes can query the oracle to ask who the ISP prefers a particular node neighbor with. Since the ISP can choose its preferences on any basis it could also facilitate traffic attraction.

Rewriting packets: In P2P systems, nodes will often ask other nodes or trackers about who else they know in the overlay. ISPs can modify the packets that are sent in response to these requests so that the asking node will neighbor with the nodes of the ISP's choosing. An example of this is shown in Figure 2.

Invisible proxy: Instead of modifying the packet, an ISP can set up a proxy that will return a response with the nodes of its choosing, acting as if it is the intended recipient.

Ono: If a node's CDN redirects are similar to another node, it will preferentially exchange data with it. Thus a node that wants to attract traffic can simply mimic another node's redirects back to it, tricking the node into thinking that they are actually close by.

IV. Evaluation Results

We aim through simulations driven by real-world data to determine the effectiveness of traffic attraction. We consider the scenario where the attractor is a transit Autonomous System (AS). In this work, we use ISP and AS interchangeably. The attractor's aim is to increase the volume of traffic flowing through it on profitable customer paths. This type of traffic attraction is the transit attraction as explained in Section III. We first describe our datasets and then present our results.

A. Datasets

To get a realistic view of how traffic attraction would affect P2P and CDN traffic we obtain real data from BitTorrent crawls and find the IP addresses of Akamai's and YouTube's world-wide deployment of replicas. We also employ BGP data to understand how traffic flows across transit ASes.

1) P2P Crawls: We crawled the popular tracker OpenBitTorrent [8] every hour for 8 days. Every torrent that had at least one active downloader was crawled. As many users are behind NATs, the tracker considers a unique (IP, port) combination as a peer. We then map peers to ASes by using the Team Cymru service [9]. Over the 8 days we found 138 million peers in over 2.75 million torrents. These peers are in over 12,000 ASes.

2) CDN Crawls: To find CDN replicas we employ a methodology similar to Huang et al. [10]. As CDNs redirect clients based on their location, we first find open RDNS servers all over the world of which we were able to find over 60,000 in 8,638 ASes. Using these RNDS servers as a platform, we query them with many different Akamai and YouTube owned domain names. We were able to find over 51,600 different Akamai servers in 578 ASes. We were also able to find 3,300 YouTube servers in 88 ASes. While we do not have means to verify that our datasets cover the complete Akamai or YouTube deployment, we note that the effectiveness of traffic attraction can only *improve* as the number of known CDN replicas increases.

3) Inter-AS Topology and Routing: Since we aim to know the effectiveness of traffic attraction on the Internet today, we must know not only in which ASes do nodes in the overlays reside, but we must also infer the transit ASes that carry the traffic between these ASes. We use the algorithm proposed by Qiu et al. [11] to find these paths. First, the algorithm determines the business relationships between ASes with Gao's inference algorithm [12], then it determines the most likely valley-free path between ASes. We use these paths and business relationships to calculate how much traffic ASes pass to their providers, customers and peers. As input for Qiu et al.'s algorithm we use Routing Information Bases (RIBs) provided by the Oregon's RouteViews Project [13]. This set of routing table dumps represents over 329,000 prefixes from 33,910 ASes.

B. BitTorrent Results

In BitTorrent, the tracker plays a crucial role in peer discovery and ultimately decides the pool of nodes that one downloads from. Therefore, we consider the scenario where a transit AS controls the tracker and can choose nodes intelligently to increase its traffic. This gives us the ability to understand the full potential of traffic attraction. We focus our study on the last day of BitTorrent crawls obtained which allows us to know who already involved in the torrents is a seeder to the rest of the participants. Given this set-up, we then simulate how much traffic flows through the transit AS on customer paths. We evaluate the increase in customer traffic by simply dividing the amount of traffic on customer links after traffic attraction by the amount of traffic on customer links originally. We determine how much traffic is on links originally by simulating when nodes are selected randomly by the tracker.

We found 1064 transit ASes that can potentially do traffic attraction and show our results for them. We plot the results when transit ASes do regular traffic attraction in Figure 3(a).

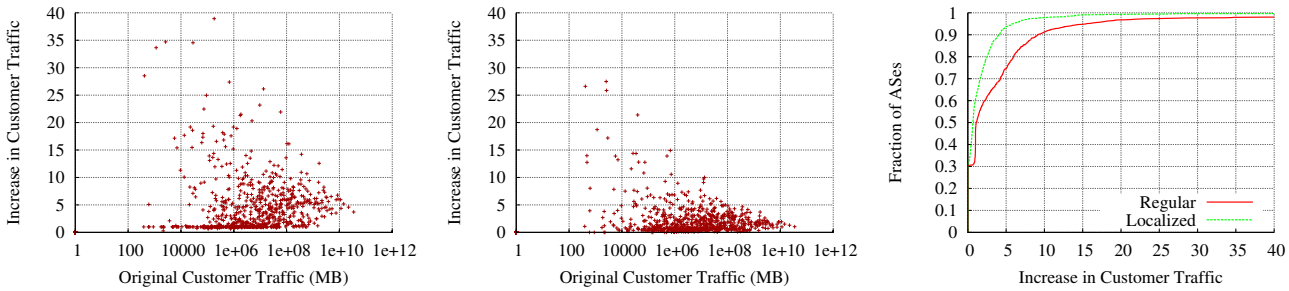

(a) Regular traffic attraction where each point represents when a single AS does the attraction.

(b) Localized traffic attraction where each point represents when a single AS does the attraction.

(c) CDF of regular and localized transit traffic attraction.

Fig. 3. Increase in customer traffic under different scenarios.

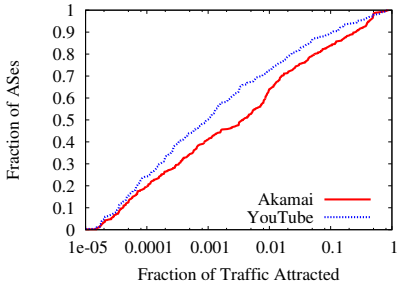

Fig. 4. Potential amount of CDN traffic that a transit AS can attract.

Many of the largest ASes are able to double the traffic on customer links. This would undoubtedly increase their revenues greatly. Smaller transit ASes can increase their customer traffic even more though, with some being able to increase them over 20 times. We also evaluate the increase in customer traffic when, instead of forcing a residential AS to download as many copies as the number of peers involved in the torrent, the transit AS is "nice" and requires only the necessary one copy of the file to be downloaded. We show the results in Figure 3(b) and compare the two scenarios by plotting two CDFs in in Figure 3(c). This shows that even if the transit attacker allows localization of P2P traffic, and does not increase the amount of money residential ASes must pay, they can still increase their profits significantly by "stealing" traffic from other transit ASes.

C. CDN Results

We now consider how much traffic transit ASes can attract from CDNs. As we do not have a baseline by which to know how much CDN traffic flows currently through a particular transit AS, we therefore cannot calculate how much traffic attraction will increase this traffic. We instead assume that traffic from CDN replicas to residential ASes is proportional to the population of end-hosts in that AS. We estimate this population with the numbers gathered from our BitTorrent crawls. We can then measure the fraction of the total CDN traffic that a single transit AS can potentially attract to go through customer links.

We calculate the results from the Akamai and YouTube datasets and plot both in Figure 4. We show results only for the 520 transit ASes that can attract Akamai traffic and 280 transit ASes that can attract YouTube traffic. From the figures, we see that most ASes only have the opportunity to carry a small fraction of CDN traffic. However, we suspect that since CDNs normally provide good locality that much of this traffic will be extra traffic that it was not seeing before and thus will result in an increase in revenue. On the other hand, we also see over 10% of ASes are able to attract over 10% of all CDN traffic. Which, given the popularity and world-wide nature of these two CDNs, is potentially a large amount of traffic.

V. Future Work and Conclusion

In this work, we investigate means by which a transit AS can increase the amount of traffic on revenue-generating links by exploiting overlays. We identify several different ways that an AS can conduct traffic attraction for both P2P systems and also CDNs. We then, through real-world datasets, evaluate how effective traffic attraction can be. Future work involves designing defenses for CDN and P2P systems, enabling them to make sure that their infrastructure can not be influenced by outside entities. We will also investigate more deeply to understand why particular ASes can more effectively attract traffic than others.

References

[1] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4P: Provider Portal for Applications," *SIGCOMM*, 2008.
[2] V. Aggarwal, A. Feldmann, and C. Scheideler, "Can ISPs and P2P Users Cooperate for Improved Performance?" *SIGCOMM Computer Communication Review*, vol. 37, no. 3, pp. 29–40, 2007.
[3] D. R. Choffnes and F. E. Bustamante, "Taming the Torrent: a Practical Approach to Reducing Cross-ISP Traffic in Peer-to-Peer Systems," in *SIGCOMM*, 2008.
[4] F. Picconi and L. Massoulié, "ISP-Friend or Foe? Making P2P Live Streaming ISP-aware," in *ICDCS*, 2009.
[5] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, and A. Zhang, "Improving Traffic Locality in BitTorrent via Biased Neighbor Selection," in *ICDCS*, 2006.
[6] B. Cohen, "Incentives build robustness in BitTorrent," in *Proc. of P2P Economics*, 2003.
[7] R. Perdisci, M. Antonakakis, X. Luo, and W. Lee, "WSEC DNS: Protecting recurse DNS resolvers from poisoning attacks," in *DSN*, 2009.
[8] OpenBittorrent, http://www.openbittorrent.com.
[9] Team-Cymru, "IP to ASN mapping," http://www.team-cymru.org/Services/ip-to-asn.html.
[10] C. Huang, A. Wang, J. Li, and K. W. Ross, "Understanding hybrid cdn-p2p: why limelight needs its own red swoosh," in *NOSSDAV*, 2008.
[11] J. Qiu and L. Gao, "AS Path Inference by Exploiting Known AS Paths," in *GLOBECOM*, 2004.
[12] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 733–745, 2001.
[13] O. RouteViews, http://www.routeviews.org/.