

# Economics of Traffic Attraction by Transit Providers

Pradeep Bangerla  
Carlos III University of Madrid, Spain

Sergey Gorinsky  
IMDEA Networks Institute, Spain

**Abstract**—Internet transit is an economy where providers sell traffic-delivery services. Traffic attraction refers to BGP (Border Gateway Protocol) techniques enabling an AS (Autonomous System) to receive traffic that would otherwise flow elsewhere. Unlike prior studies that present security perspectives on traffic attraction or deal with economic considerations in game-theoretic settings, this paper focuses on the economics of customer-traffic attraction by transit providers and report extensive simulations in an Internet-scale model configured with realistic data on traffic, topology, and pricing. We consider traffic attraction by tier-1, tier-2, and tier-3 networks with 3 types of reactions by other networks: filtering, customer-disconnections, and attempts of losing ASes to attract traffic to themselves. Our results demonstrate that transit providers can derive substantial financial benefits from attracting customer traffic, with tier-1 networks being in the strongest position to do so. The traffic attraction remains effective despite the countermeasures unless participation in them is broad.

**Keywords**- BGP; economics; transit provider; traffic-delivery payment; traffic attraction; incentive; filtering; disconnection

## I. INTRODUCTION

Internet transit is a massive decentralized economy where thousands of providers sell and resell traffic-delivery services. Transit revenues depend on traffic rates: the more traffic the provider transits, the larger the revenue is. Persistent declines in transit prices have recently put transit businesses under significant financial pressure. Furthermore, peering [1], [2], caching [3], joint purchase of transit in bulk [4], [5], and other cost-reduction techniques by customers tend to decrease revenues of transit providers.

Traffic attraction refers to a family of BGP (Border Gateway Protocol) [6] techniques enabling an AS (Autonomous System) to receive traffic that would otherwise flow elsewhere. While some of prior studies on traffic attraction [7]–[16] indicate that attraction of extra traffic provides economic benefits to transit providers, the prior work mostly focuses on security rather than traffic economics. The previous security-themed work examines various traffic-attraction techniques and argues that secure versions of BGP or additional router mechanisms do not neutralize them. There is also a body of related game-theoretic and simulation-based studies that include economic considerations but analyze traffic attraction in small-scale artificial settings.

In this paper, we focus on the economics of customer-traffic attraction by transit providers and report extensive C-BGP [17] simulations in an Internet-scale model configured with realistic data on inter-domain traffic, topology, and

pricing. We consider attractors from the top 3 tiers of the transit hierarchy as well as 3 types of reactions by other ASes to the attraction: (1) filtering, i.e., discarding the BGP announcements that trigger the attraction, (2) disconnection by discontented customers, i.e., severance of their business relationships with the attractor altogether, and (3) attempts of discontented ASes to attract extra traffic to themselves. The broader scope and higher realism – combined with sensitivity studies – enable our work to offer deeper quantitative insights into traffic-attraction economics and reach reliable qualitative conclusions.

The results demonstrate that tier-1, tier-2, and tier-3 networks have significant financial incentives to attract traffic. The financial gains from attraction is derived by pulling extra traffic from the peering links up the transit hierarchy. Our paper shows that the tier-1 ASes are in the strongest position to benefit from traffic attraction and preserve their own gains when multiple networks attempt to attract traffic. The traffic attraction remains effective despite filtering and disconnections by other ASes unless the participation in the countermeasures is broad. Filtering is ineffective even when it involves all the ASes that financially suffer from the attraction. The disconnection by discontented customers becomes effective only when a large portion (e.g., a half) of them terminate their business relationships with the attractor. While the traffic attraction slightly increases router complexity and elongates AS-level paths, these effects are too insignificant to be strong deterrents against the attraction. Our studies of the model sensitivity to the topology, traffic, and pricing show that the quantitative outcomes are less sensitive to the traffic matrix than to the topology and pricing. Our results provide a foundation for the future studies on the economics of traffic attraction by transit providers.

## II. ATTRACTION VIA PREFIX DEAGGREGATION

The specific BGP technique for traffic attraction in our study is prefix deaggregation. Prefix deaggregation refers to representing a prefix with multiple longer prefixes and announcing these longer prefixes to other ASes via BGP. Due to the longest-prefix match rule of IP forwarding, the BGP announcement of a longer prefix steers traffic to the announced path. Multihomed ASes routinely employ prefix deaggregation to balance their incoming traffic among their multiple connections to the Internet. Instead of the load balancing by multihomed ASes, our paper studies attraction of additional customer traffic. We consider the kind of prefix deaggregation




Fig. 1. Example of traffic attraction via prefix deaggregation: (a) traffic routing before the attraction; (b) traffic routing after the attraction.

where an intermediary AS learns a prefix from a customer, deaggregates the prefix, and announces all longer prefixes to each of its other customers. In particular, the traffic-attracting AS splits a learned prefix equally into 2 longer prefixes and announces both longer prefixes to the customers (Note that the AS deliberately does not announce the deaggregated prefixes to any of its peers so that none of its current traffic shifts from its customer links to its peering links).

To illustrate how prefix deaggregation enables traffic attraction, figure 1 presents a simple example where networks *B* and *C* directly learn prefix *x* from their mutual customer *A*. When providers *B* and *C* propagate prefix *x* to their another mutual customer *D*, network *D* sends traffic to network *A* through provider *B*, and network *B* collects transit payments from both customers *A* and *D*. When provider *C* deaggregates prefix *x* into longer prefixes *y* and *z*, and announces both deaggregated prefixes *y* and *z* to customer *D*, the traffic from network *D* to network *A* is attracted to flow through provider *C* rather than provider *B*; consequently, the traffic-delivery payments from both customers *A* and *D* go to provider *C* instead of provider *B*.

While the considered prefix-deaggregation method is easily implementable in practice, a transit AS can also attract extra traffic by employing a different BGP technique. For example, the intermediate AS can attract additional traffic by reducing the path length in a propagated BGP announcement. In practice, prominent ASes attract extra customer traffic by rewriting the origin attribute in propagated BGP announcements [18], [19]. These and other alternative techniques for traffic attraction represent an interesting topic for future studies on traffic economics.

For the role of traffic attractors, we select ASes throughout the Internet transit hierarchy. Whereas tier-1 networks [20] are in the strongest position to attract significant amounts of extra traffic, our interactions with industry experts suggest that traffic attraction by tier-2 ASes is more common. In addition to traffic attraction by tier-1 and tier-2 transit providers, we also examine traffic attraction by tier-3 networks.




Fig. 2. Distribution of origin traffic shares in our traffic matrices.

### III. MODEL

Our modeling strives to abstract the highly complicated problem into a manageable representation for realistic study and discussion. Instead of focusing on a single setting, we parameterize our AS-level Internet model to experiment with realistic ranges of parameter settings.

#### A. Topology

Despite a decade of intensive research, the AS-level Internet topology is not known accurately, e.g., due to missed links [21]. To deal with this uncertainty, we consider 3 alternatives including topologies reported by CAIDA [22] on 2012/6/1 and UCLA [23] on 2012/5/1. The third topology, to which we refer as UCLA+, is synthetic. We derive it from the UCLA topology by adding links that UCLA reported for at least 15 out of the 30 subsequent days. The enhancement contributes 1,656 peering and 75 transit links to the UCLA+ topology.

#### B. Traffic

Characterizing the inter-domain traffic is another notoriously hard problem. Our study considers 9 traffic matrices guided by empirical Internet data. The matrices generalize the measurement results suggesting that (a) the fraction of traffic originated by the largest content source was about 5% in 2009 and growing [24], and (b) the distribution of traffic from origin ASes is Zipf-like with the shape parameter between 0.9 and 1.1 [24], [25]. We obtain the  $9 = 3 * 3$  traffic matrices by combining 3 settings along each of the following 2 dimensions: (1) overall distribution function of origin traffic shares and (2) assignment of the traffic shares to specific ASes.

For the overall distribution of the origin traffic shares, we consider 3 instances of the Zipf-Mandelbrot function defined as  $K/(r+v)^D$ , where  $r$  refers to the rank of the AS,  $D$  denotes the shape parameter, and  $K$  and  $v$  are constants. The instance settings are as follows:

- [Z1]  $D = 1.1$ ,  $K = 0.181$ , and  $v = 0.5$ ;
- [Z2]  $D = 0.9$ ,  $K = 0.072$ , and  $v = 0$ ;
- [Z3]  $D = 0.7$ ,  $K = 0.108$ , and  $v = 0$ .

Figure 2 depicts the 3 overall traffic-share functions that cover the realistic ranges for the shape parameter (from 0.7 to 1.1) and largest traffic share (from 3% to 12%).

To assign the traffic shares to specific ASes, i.e., to set rank  $r$  of each AS, we consider the following 3 diverse options for the ranking metric:

- [R1] peering coefficient, which is the number of peers divided by the total number of providers and customers of the AS; the peering coefficient tends to be the highest among small/intermediate ASes that provide access to content hosts (these ASes have many peers but few transit links);
- [R2] IP address count, i.e., the number of IP addresses originated by the AS;
- [R3] website visit frequency, i.e., the number of users who visit the websites hosted by the AS.

We use notation  $Z\gamma R\delta$  to denote the traffic matrix with overall traffic-share function  $Z\gamma$  and AS ranking metric  $R\delta$ .

In each of the 9 traffic matrices, we distribute the originated traffic of an AS among destinations (i.e., content-consuming ASes) in proportion to the IP address counts of the destination ASes. In addition to the above content traffic, our traffic matrices also incorporate peer-to-peer traffic, with the peer-to-peer traffic between a pair of ASes set proportionally to the IP address count of each AS in the pair.

### C. Pricing

In comparison to traffic, pricing of inter-AS relationships is even more difficult to infer due to confidential bilateral agreements guarding them. Empirical evidence indicates that IP transit is subject to subadditive pricing where prices per Mbps are lower for higher traffic rates. In particular, the empirical data suggest as reasonable the following pricing function [1], [2], [26]: the monthly payment from the customer to the provider of transit link  $l$  is  $t_l = b \cdot L^m$  where  $L$  denotes the traffic rate in Mbps,  $b = 45$ , and  $m = 0.75$  (the computed payment is in U.S. dollars). Our study generalizes this transit pricing function by considering the range of  $m$  values from 0.4 to 1, i.e., from heavy economies of scale to linear pricing which does not offer any price discount for higher traffic rates.

While the 95th-percentile billing [27], [28] is common in IP transit charging, this method requires a large number of individual traffic-rate samples. Instead of modeling the traffic at a fine-grained time scale, we directly represent the traffic of each transit link as the 95th-percentile rate because: (a) fine-grained traffic statistics are not publicly available for most ASes, and (b) simulations for a detailed traffic model do not scale to a large population of ASes.

Peering relationships are usually free of financial settlements between the peers. On the other hand, maintaining a peering link does involve costs, e.g., payments to an IXP (Internet eXchange Point) [21] that provides the physical infrastructure for peering. To represent the peer’s cost of maintaining its peering link  $l$ , we adopt peering-cost function  $p_l = 20 \cdot L^{0.4}$  suggested in prior work [1], [2], [26]. In comparison to the transit-link payment for the same traffic rate, the peering cost is always lower but not negligible.

To determine the overall traffic-delivery payment for each AS, we partition all external links of the AS into 3 sets: set  $\mathbb{V}$  contains the transit links where the AS is a provider, set  $\mathbb{U}$

is for the transit links where the AS acts as a customer, and set  $\mathbb{G}$  includes all its peering links. Then, the monthly traffic-delivery payment of the AS is  $P = \sum_{l \in \mathbb{V}} t_l - \sum_{l \in \mathbb{U}} t_l - \sum_{l \in \mathbb{G}} p_l$  with positive values denoting traffic-delivery revenues, and negative values representing traffic-delivery expenses.

## IV. EVALUATION METHODOLOGY

To evaluate traffic attraction and countermeasures by other ASes, we conduct large-scale simulations in C-BGP [17]. C-BGP determines AS-level paths for all traffic and bidirectional traffic rates for every inter-domain link in the simulated topology. The BGP routing policies in C-BGP satisfy the valley-free routing conditions [29]. Although real Internet routing occasionally deviates from such policies, valley-free routing constitutes a reasonable approximation because we are mostly interested in qualitative insights.

We optimize C-BGP to overcome its scalability limitations. For example, even if each AS announces a single prefix only, the standard C-BGP can exhaust a relatively large physical memory before computing stable AS-level paths. We improve the memory management of C-BGP to scale up the simulations to at least 6,000 prefix announcements and supply the optimized version to the online C-BGP code repository.

While our optimizations alleviate – but not eliminate – the C-BGP scalability limitations, we conduct the simulations by focusing on the core of each topology and representative prefixes of the ASes. We extract the topological core by excluding all the stub ASes and all their links. For the 3 examined topologies, their cores contain around 6,000 transit ASes. We determine traffic-delivery payments for all these transit ASes. Although the simulations do not directly consider the stub ASes outside the topological core, we account for the traffic of these ASes when computing the traffic-delivery payments for the transit ASes. While real ASes generally own and announce several prefixes, the C-BGP scalability limitations prevent announcing multiple prefixes from each AS. On the other hand, a single prefix per AS is sufficient for C-BGP to simulate all inter-domain communications in the topological core. Therefore, we associate the total inter-domain traffic of each AS with a single representative prefix of this AS.

To configure our traffic matrices, we rely on real data. Based on the Cisco-VNI statistics [30], we set the total rates of the inter-domain content and peer-to-peer traffic to 45 Tbps and 15 Tbps respectively, with 4 times more content traffic flowing from servers to clients than from clients to servers. We specify the IP address counts of the ASes according to the CIDR report [31]. The peering-coefficient ranks of ASes are in accordance with the CAIDA topology. We calculate the website visit frequencies based on the data provided by the Alexa Web Information Service from Amazon for top 100,000 websites [32].

Among the 6,000+ transit ASes in the topological cores, we select 30 ASes to act as traffic attractors. Each of tiers 1, 2, and 3 in the transit hierarchy contributes exactly 10 ASes (those with the largest numbers of customers) to the attractor




Fig. 3. Payment change for the attracting AS.




Fig. 4. Payment changes for the tier-1 attractors.

set. We refer to the 10-AS groups as T1, T2, and T3 respectively. To denote AS  $\beta$  from tier  $\alpha$ , we use notation  $T\alpha\beta$ .

Our main metric is payment change. It measures the relative change in the traffic-delivery payment of the AS in comparison to the baseline scenario where no AS tries to attract additional traffic.

Unless explicitly stated otherwise, we report the results for the following default settings. The topology is from CAIDA. The results are averaged over the 9 traffic matrices. Transit-pricing exponent  $m$  is set to 0.75. When deaggregating prefixes to attract traffic, the attracting AS deaggregates prefixes announced by its 100 largest customers. We also study sensitivity of the results to the topology, traffic, and pricing.

## V. EVALUATION RESULTS

### A. Attraction by a single AS

We start by examining what happens when a single AS attempts to attract traffic. We repeat this experiment for the 30 attractors with each of the 9 traffic matrices and record the payment change for the attracting AS. Using box plots, figure 3 presents the results arranged according to the tier of the attracting AS. The plots demonstrate that transit ASes have significant financial incentives to attract traffic: the median payment change is 148%, 38%, and 21% for T1, T2, and T3 respectively. The tier-1 networks are in the strongest position to benefit from traffic attraction because they sell transit to numerous customers but do not buy transit themselves. Being less central in the transit hierarchy, all considered tier-2 ASes




Fig. 5. Payment changes for the winning and losing ASes when the attracting AS is from tier 1, 2, or 3.

are still able to raise their revenues by attracting traffic. While some tier-3 networks also benefit substantially from the traffic attraction, it is not beneficial for other attracting ASes from T3.

Focusing on the experiments with the tier-1 attractors, figure 4 shows that while T1j and T1c are 2 opposite extremes with the median payment changes of 830% and 28% respectively, the payment changes for T1b and the 7 other ASes from T1 are rather similar to each other. The large gap between the payment changes for T1j and T1c is mostly due to the different sizes of these ASes. In comparison to T1j, T1c serves more transit traffic and attracts a larger amount of extra traffic in absolute terms. Nevertheless in relative terms, the payment gain is much higher for the smaller T1j.

### B. Winners, losers, and neutrals

By redistributing traffic in the AS-level topology, traffic attraction by an AS affects traffic-delivery payments for other ASes. We refer to the ASes with increased traffic-delivery payments as winners, and to the ASes with negative payment changes as losers. Neutrals are the ASes with unchanged traffic-delivery payments. Figure 5 plots the payment changes for the winning and losing ASes in the experiments of section V-A. The traffic attraction by tier-1 networks makes the most divisive impact on the payments: the fractions of winners, losers, and neutrals are 17%, 26%, and 57% respectively. When the attracting AS is from T2, the fractions of winners and losers decrease to 11% and 8% respectively. After the traffic attraction by tier-3 networks, the impact is highly local: either winners or losers comprise only 1% of the AS population.

### C. Impact on traffic

To understand where the attracted traffic comes from, we classify the inter-domain traffic of an AS into 3 types: (1) upstream, i.e., traffic from the AS to its providers, (2) downstream, i.e., traffic from the AS to its customers, and (3) peering, i.e., traffic on the peering links of the AS. Then, we consider the 10 ASes that suffer the largest declines (in absolute terms) of their traffic-delivery payments after the traffic attraction by the T1 networks. These largest losers are 4 tier-1, 5 tier-2, and 1 tier-3 ASes. Figure 6a depicts changes in the upstream, downstream, and peering traffic of

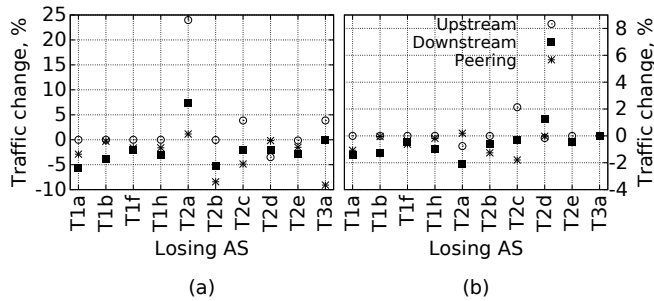


Fig. 6. Changes in the upstream, downstream, and peering traffic of the 10 largest losers when the attractor is a tier-1 AS: (a) without filtering; (b) with filtering.

these 10 losers after the attraction. While the losing tier-1 ASes have no upstream traffic either with or without the attraction, their downstream traffic decreases, suggesting that the lost downstream traffic is acquired by the attractor. The attraction makes a qualitatively similar impact on the traffic of T2b and T2e. For the 3 other largest losers from the lower tiers, the upstream traffic increases due to sending more traffic to the tier-1 attractor, and the downstream traffic increases to a smaller extent (e.g., for T2a) or even decreases (e.g., for T2c). From an overall perspective, the traffic attraction reduces peering traffic and pulls extra traffic up the transit hierarchy.

#### D. Filtering by losing ASes

To analyze responses of other ASes to the traffic attraction, we first consider filtering, i.e., discarding the deaggregated prefixes announced by the attractor. Figure 7 presents the payment change for the attracting AS when all losing ASes from the experiments in section V-A do the filtering. Comparing the results in figures 3 and 7, we see that the filtering reduces but does not remove the financial benefits for the traffic attractor. With the filtering, the median payment change for the attracting AS is 37%, 11%, and 4% for T1, T2, and T3 respectively.

Figure 6b offers insights into the inability of the filtering to negate the attraction. Although the filtering can help a losing AS – e.g., T2a – to reduce its upstream traffic to the attractor, the filtering by the loser does not prevent its customers from their switching to alternative paths via the attractor.

#### E. Multi-stage rational filtering

The filtering examined in section V-D redistributes traffic in the AS-level topology and creates additional losers. In this section, we extend the above filtering scenario into a multi-stage reaction where the group of filtering ASes on each stage expands by incorporating the additional losers from the previous stage. We refer to the filtering by losers as rational because it is done only by the ASes that financially suffer from the traffic redistribution.

To assess the effectiveness of multi-stage rational filtering, we consider the setting where stage 0 corresponds to the traffic attraction by T1b without filtering. Figure 8 shows that while




Fig. 7. Payment change for the attracting AS when all losers do the filtering.




Fig. 8. Payment change for attractor T1b with multi-stage filtering.

the payment change for attractor T1b decreases on stages 1 through 3, the payment gain for the attractor stabilizes at 4% after stage 3, which yields no additional losers. Hence, even if the filtering is done by all losing ASes, the multi-stage rational filtering does not eliminate financial incentives for the traffic attraction. The use and propagation of the deaggregated prefixes by the winners and neutrals allow the attractor to increase its traffic-delivery revenues.

#### F. Cooperative filtering

While section V-E demonstrates the inability of rational filtering to negate the financial benefits of the attractor, we now explore what happens if the other winning ASes go against their own financial interests and also react by filtering the deaggregated prefixes. Referring to such filtering as cooperative, we consider the multi-stage cooperative filtering where stage 0 corresponds to the traffic attraction by T1b without filtering, stage 1 involves filtering by all losers and neutrals, and each subsequent stage expands the group of the filtering ASes with additional winners (selected in the increasing order of their payment gain). Figure 8 shows that the cooperative filtering negates the payment gain of attractor T1b on stage 5 where the filtering is done by all customers of the attractor.

#### G. Disconnection

The evaluation in sections V-D through V-F shows that filtering is not an effective countermeasure unless the winners resist the traffic attraction against their own financial interests. Now, we examine a more severe reaction by losers where




Fig. 9. Payment change for attractor T1b with multi-stage disconnection by losing customers of the attractor.




Fig. 10. Payment changes for 10 ASes when the traffic attraction is done by T1b only vs. all 10 ASes.

losing customers sever their business relationships with the attracting AS altogether. Again, we consider a multi-stage version of the response where stage 0 corresponds to the traffic attraction by T1b without filtering (and without disconnection). On stage 1 of the disconnection response, the attractor is disconnected from the 1% of its losing stage-0 customers that are selected in the decreasing order of their absolute losses. On each of stages 2 through 7, the cumulative number of the disconnected customers of T1b doubles. On stage 7 where the attractor is disconnected from 45% of all its original customers (i.e., 64% of its losing stage-0 customers), the attractor still has the payment gain of 8%. On each of the subsequent stages, we disconnect all remaining T1b losing customers from the previous stage. Both stages 8 and 9 create additional losers among the connected T1b customers. On stage 10 where the attractor is disconnected from 85% of all its original customers, no new losers emerge, and the payment change of T1b stabilizes. The remaining 15% of all original T1b customers are either winners or neutrals on stage 10 and hence do not disconnect from the attractor. Figure 9 depicts the dynamics of the multi-stage disconnection response by the losing customers. The results demonstrate that the disconnection by losing customers is ineffective unless a large portion of them terminate their business relationships with the attractor.




Fig. 11. Changes in the upstream, downstream, and peering traffic of 10 ASes when the attraction is done by: (a) T1b only and (b) all 10 ASes.

### H. Attraction by multiple ASes

The previous sections show that neither filtering nor disconnection eliminates the financial incentives for traffic attraction unless participation in the response is broad. Now, we consider a different reaction where a losing AS defends its traffic-delivery payment by attracting extra traffic to itself. Specifically, we consider the scenario where in response to the traffic attraction by T1b the 9 largest losers (in absolute terms) try to attract traffic as well. The expanded set of 10 attractors includes 6 tier-1, 3 tier-2, and 1 tier-3 ASes. Figure 10 shows the payment changes for these 10 ASes when the traffic attraction is done by T1b only vs. all 10 ASes. When all 10 ASes try to attract traffic, all 6 tier-1 ASes and T2f gain from the traffic attraction but the payment changes for T2b, T2g, and T3a are negative. The results confirm our earlier observation that tier-1 networks are in the strongest position to benefit from traffic attraction. Furthermore, figure 10 demonstrates that ASes from lower tiers are not assured to gain from traffic attraction when multiple networks attempt to attract extra traffic.

To understand why the transit ASes from different tiers fare differently when multiple networks try to attract traffic, figure 11 plots the changes in the upstream, downstream, and peering traffic of the 10 ASes. When T1b acts as the only attractor, it greatly increases its own downstream traffic and decreases the downstream traffic for the other ASes except T3a that suffers the loss due to an increase in the upstream traffic. When all 10 ASes try to attract traffic, all 6 tier-1 networks (which never have any upstream traffic) win by increasing their downstream traffic. For T2f, the downstream traffic grows as well, and this growth outweighs the increase in its upstream traffic. On the other hand, T2b, T2g and T3a remain losers because their upstream traffic grows significantly while their downstream traffic increases less (if at all). Some attractors are more powerful than others.

### I. Sensitivity to the topology

To study how sensitive our model is to its assumptions, we first examine its sensitivity to the AS-level Internet topology in the scenario where only T1b attracts traffic. In addition to the CAIDA topology, we also consider the UCLA and UCLA+ topologies. Figure 12 presents the payment changes




Fig. 12. Sensitivity of payment changes to the AS-level Internet topology.

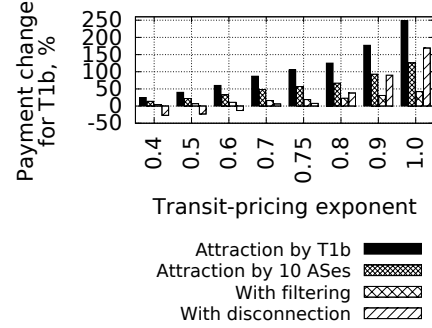


Fig. 14. Sensitivity of the attractor's payment change to transit pricing.




Fig. 13. Sensitivity of the attractor's payment change to the traffic matrix.

for the attractor (which is the largest winner) and 5 largest losers. In quantitative terms, the topology has a substantial impact, e.g., the payment change for the attracting AS varies from 85% to 180% with the CAIDA and UCLA+ topologies respectively. On the losing side, the maximum loss by an AS varies from -9% to -81% with UCLA and UCLA+ topologies respectively. This latter result highlights the importance of the topology in general and peering links in particular for economic outcomes of traffic attraction. By adding the relatively small numbers of 75 transit and 1,656 peering links, the enhancement of the topology from UCLA to UCLA+ makes the traffic attraction significantly more powerful. The distributions of payment changes for all ASes are qualitatively the same in the 3 examined topologies.

#### J. Sensitivity to the traffic matrix

Expanding the above sensitivity study, we now examine the role of the traffic matrix. For each of our 9 traffic matrices, figure 13 plots the payment change for attractor T1b in 4 scenarios: (1) only T1b attracts traffic, (2) attraction is by 10 networks as in section V-H, (3) filtering is done by all losing networks as in section V-D, and (4) 50% of the T1b losing customers disconnect from the attractor. The 9 considered traffic matrices are diverse with respect to both overall distribution of origin traffic shares and assignment of the traffic shares to specific ASes. Despite this diversity, the results are qualitatively the same and relatively stable: the payment change for T1b varies from 75% to 105% when only

T1b attracts traffic, from 34% to 57% with the 10 attractors, from 18% to 27% with the filtering, and from 7% to 26% with the disconnection. The quantitative outcomes of the traffic attraction are less sensitive to the traffic matrix than to the topology.

#### K. Sensitivity to pricing

To assess the sensitivity of our model to pricing, we consider the same 4 scenarios as in section V-J and reduce transit-pricing exponent  $m$  from 1 to 0.9, 0.8, 0.75, 0.7, 0.6, 0.5, and finally 0.4. Figure 14 tracks the payment change for attractor T1b and exhibits substantial quantitative variations in the outcomes. In the first 3 scenarios, the payment change for T1b remains positive but decreases greatly: from 249% to 25% when T1b is the only attractor, from 126% to 14% when the 10 ASes attract traffic, and from 42% to 4% with the filtering. With the disconnection by 50% of the losing customers, the payment change not only decreases as the transit-pricing exponent is reduced but also becomes negative: the payment change for T1b is 169% for  $m = 1$  and -27% for  $m = 0.4$ . The payment change becomes 0 when the transit-pricing exponent is around 0.65. According to Telegeography data, the transit-pricing exponent is currently 0.8, 0.73, 0.69, and 0.65 for Oceania, Asia and South America, Europe, and North America respectively. Thus, our results suggest that at least a half of the losing customers needs to disconnect from the attractor to eliminate the benefits of the attraction under the current pricing.

#### L. Sensitivity to attraction intensity

While sections V-E through V-G explore how intensive the filtering and disconnection should be to negate the benefits of traffic attraction, we now examine the sensitivity of the attractor's gain to the intensity of attraction. In these experiments where only T1b attracts traffic (as in section V-A), T1b changes its attraction intensity by deaggregating a different number of prefixes announced by its largest customers. Specifically, the attractor deaggregates 1, 10, 100, or 1,000 prefixes. Figure 15 presents the payment change for T1b. The median payment change is 7%, 30%, 78%, and 108% for 1, 10, 100, and 1000 deaggregated prefixes respectively. The marginal utility of the attraction intensity diminishes quickly.




Fig. 15. Sensitivity of the attractor’s payment change to attraction intensity.

Even by deaggregating a relatively small number of prefixes (e.g., 100 as in the default setting of our studies) the attractor obtains most of its maximum possible gain.

#### M. Impact on router complexity

Traffic attraction via prefix deaggregation makes network routers more complex due to the need to store and process the deaggregated prefixes. While section V-L demonstrates that the attractor can obtain most of its benefits by deaggregating a much smaller number of prefixes than the routers handle currently, the respective increase in router storage and processing is insignificant. Based on the costs of route-processor and line cards for modern Juniper routers, we analyze the amortized monthly cost of storing the extra prefixes (details of the analysis are omitted due to space constraints) and conclude that it is multiple decimal magnitudes lower than the traffic-delivery payment gains from the attraction. This conclusion is consistent with prior game-theoretic analyses of incentives for prefix deaggregation [11]. Therefore, the increased router complexity is not a strong deterrent against the traffic attraction.

#### N. Impact on path lengths

Finally, we evaluate how the traffic attraction and countermeasures affect the lengths of AS-level paths. Figure 16 plots the distributions of the AS-level path lengths for 5 scenarios: (1) baseline without any attraction; (2) attraction by T1b only, (3) attraction by 10 networks as in section V-H, (4) with stabilized multi-stage rational filtering by all losing networks as in section V-E, and (5) with stage-7 disconnection by losing customers of T1b as in section V-G. The path-length distributions are mostly similar to each other. The average path length is 4.4 hops in the baseline scenario. When only T1b attracts traffic, paths elongate, with their average length increasing to 5.1 hops. The attraction by the 10 ASes reduces this elongation, with the average path length becoming 4.8 hops. The filtering almost restores the path-length distribution of the baseline scenario and yields the average path length of 4.5 hops. The disconnection has an opposite effect, with the average path length increasing to 5.7 hops. Overall, the elongation of paths under traffic attraction is not significant to deter this revenue-increasing behavior.




Fig. 16. Distributions of AS-level path lengths.

## VI. RELATED WORK

Unlike our economic investigation of traffic attraction, prior studies of the subject approach it mostly from security perspectives [7]–[9]. Ballani et al. [7] explore the ability of an AS to attract traffic to itself for either discarding the attracted traffic or delivering the traffic to the destination. The paper considers 2 different attraction techniques where the attracting AS announces an invalid path to a prefix: by spuriously claiming to be either the prefix owner or an intermediary on the invalid path. [7] estimates the feasibility of such traffic attractions but does not study their economic impact. Nordstrom and Dovrolis [9] explore various attraction techniques and 2 countermeasures: filtering and adoption of S-BGP, a secure version of BGP. Again with an exclusive focus on security rather than traffic economics, [9] concludes that filtering is ineffective and that S-BGP is too heavy to get deployed. Goldberg et al. [8] study robustness of S-BGP and other secure routing protocols to traffic attraction. [8] argues that the secure routing protocols fail to neutralize traffic attraction and need to be supplemented with defensive filtering.

There is also related work that includes economic considerations [10]–[15]. Gill et al. [12] analyze a game where the financial benefits of traffic attraction serve as incentives to deploy S-BGP. [12] does not evaluate the economic incentives but simply use them as a basis for the analyzed S-BGP deployment scenario. Lutu et al. [15] study whether traffic engineering via prefix deaggregation can reduce the transit expenses of prefix owners. In contrast, our work evaluates traffic attraction by intermediary ASes that seek to increase their transit revenues. While Bangera and Gorinsky [14] simulate the YouTube prefix-hijacking incident to assess its economic implications, our investigation considers multiple attractors from different transit tiers and evaluates novel countermeasures, such as multi-stage filtering and disconnection. Goldberg et al. [10], Kalogiros et al. [11], and Levin et al. [13] present game-theoretic investigations of economic incentives in inter-domain routing. While these papers provide thorough analyses for small-scale artificial settings, we conduct Internet-wide simulations driven by realistic data on inter-domain traffic, topology, and pricing.



The broader scope and higher realism, combined with the sensitivity studies, enable our work to yield new insights. For example, while [7] refuses to consider deaggregation-based attraction because the attractor is presumably unable to deliver the attracted traffic to the destination, our study demonstrates that the deaggregation-based attraction and delivery are not only feasible but also highly beneficial for the transit revenues of the attractor.

## VII. CONCLUSION

Relying on extensive modeling and C-BGP simulations, this paper presents an economic perspective on traffic attraction and countermeasures. The attraction and reactions redistribute traffic in the AS-level topology and create numerous winners and losers in the AS population.

The results demonstrate that tier-1, tier-2, and tier-3 networks have significant financial incentives to attract traffic. In comparison to ASes from the lower tiers, the tier-1 networks are in a stronger position to benefit from traffic attraction with respect to: (a) the degree of the attainable gain, (b) impact on other networks, and (c) preserving their own gain when multiple ASes attract traffic. The traffic attraction provides the financial gains by pulling extra traffic from peering links up the transit hierarchy.

The traffic attraction remains effective despite countermeasures unless the participation by ASes is very broad. Rational filtering does not remove the attraction incentives even when all losing ASes do the filtering. Only if winning ASes go against their own financial interests and join the filtering, such cooperative filtering eliminates the financial benefits of the attractor. The disconnection by losing customers is ineffective unless a large portion of them terminate their business relationships with the attractor. The increased router complexity and elongation of paths are too insignificant to be strong deterrents against the traffic attraction. Our studies of the model sensitivity to the topology, traffic, and pricing show that the quantitative outcomes are less sensitive to the traffic matrix than to the topology and pricing; qualitatively, the results remain consistent.

Our paper strives to foster a discussion on traffic-attraction economics and expose deeper insights into associated trade-offs. On the one hand, the financial benefits of the attractor are substantial. On the other hand, cooperative filtering or disconnection can negate these benefits only if a very large number of ASes participates in the response. While we do not advocate (or oppose) traffic attraction, the observed trade-offs raise the possibility that the increasing financial pressure on IP transit business might prompt transit providers to attract traffic. Also, while the wide scope and sensitivity analysis make our results fairly generic, no simulation study can cover the full set of potential behaviors. Besides, this paper does not evaluate a legal perspective: while prefix deaggregation does not violate any law and is routinely used for traffic engineering, traffic attraction via prefix deaggregation might face future legal challenges. These and other extensions are interesting topics for future work.

## ACKNOWLEDGMENTS

This research was financially supported in part by the European Commission (FP7-ICT 288021, EINS) and Spanish Ministry of Science and Innovation (RYC-2009-04660).

## REFERENCES

- [1] H. Chang, S. Jamin, and W. Willinger, "To Peer or not to Peer: Modeling the Evolution of the Internet's AS-level Topology," *INFOCOM 2006*.
- [2] A. Dhamdhere and C. Dovrolis, "The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh," *CoNEXT 2010*.
- [3] S. Hasan, S. Gorinsky, C. Dovrolis, and R. Sitaraman, "Trade-offs in Optimizing the Cache Deployments of CDNs," *INFOCOM 2014*.
- [4] I. Castro, R. Stanojevic, and S. Gorinsky, "Using Tuangou to Reduce IP Transit Costs," *IEEE/ACM Transactions on Networking*, 2014.
- [5] I. Castro and S. Gorinsky, "T4P: Hybrid Interconnection for Cost Reduction," *NetEcon 2012*.
- [6] Y. Rekhter and T. Li, "A Border Gateway Protocol (BGP-4)," *RFC 1771*.
- [7] H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," *SIGCOMM 2007*.
- [8] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, "How Secure are Secure Interdomain Routing Protocols," *SIGCOMM 2010*.
- [9] O. Nordstrom and C. Dovrolis, "Beware of BGP Attacks," *ACM SIGCOMM CCR*, 2004.
- [10] S. Goldberg, S. Halevi, A. D. Jagard, V. Ramachandran, and R. N. Wright, "Rationality and Traffic Attraction: Incentives for Honest Path Announcements in BGP," *SIGCOMM 2008*.
- [11] C. Kalogiros, M. Bagnulo, and A. Kostopoulos, "Understanding Incentives for Prefix Aggregation in BGP," *ReArch 2009*.
- [12] P. Gill, M. Schapira, and S. Goldberg, "Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security," *SIGCOMM 2011*.
- [13] H. Levin, M. Schapira, and A. Zohar, "Interdomain Routing and Games," *STOC 2008*.
- [14] P. Banger and S. Gorinsky, "Impact of Prefix Hijacking on Payments of Providers," *COMSNETS 2011*.
- [15] A. Lutu, M. Bagnulo, and R. Stanojevic, "An Economic Side-Effect for Prefix Deaggregation," *NetEcon 2012*.
- [16] P. Banger and S. Gorinsky, "An Economic Perspective on Traffic Attraction by Internet Transit Providers," *ICNP 2013*.
- [17] B. Quoitin and S. Uhlig, "Modeling the Routing of an Autonomous System with C-BGP," *IEEE Network*, 2005.
- [18] The NANOG Archive, <http://mailman.nanog.org/pipermail/nanog/2012-May/048564.html>.
- [19] The NANOG Archives, <http://mailman.nanog.org/pipermail/nanog/2011-December/043036.html>.
- [20] Hasan, S. and Gorinsky, S., "Obscure Giants: Detecting the Provider-Free ASes," *Networking 2012*.
- [21] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a Large European IXP," *SIGCOMM 2012*.
- [22] CAIDA: Cooperative Association for Internet Data Analysis, <http://as-rank.caida.org/>.
- [23] Internet Research Lab, UCLA, <http://irl.cs.ucla.edu/>.
- [24] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet Inter-domain Traffic," *SIGCOMM 2010*.
- [25] H. Chang, S. Jamin, Z. M. Mao, and W. Willinger, "An Empirical Approach to Modeling Inter-AS Traffic Matrices," *IMC 2005*.
- [26] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "GENESIS: An Agent-based Model of Interdomain Network Formation, Traffic Flow and Economics," *INFOCOM 2012*.
- [27] X. Dimitropoulos, P. Hurley, A. Kind, and M. Stoecklin, "On the 95-Percentile Billing Method," *PAM 2009*.
- [28] R. Stanojevic, N. Laoutaris, and P. Rodriguez, "On Economic Heavy Hitters: Shapley value analysis of the 95th-percentile pricing," *IMC 2010*.
- [29] L. Gao and J. Rexford, "Stable Internet Routing Without Global Coordination," *IEEE/ACM Transactions on Networking*, 2001.
- [30] Cisco Visual Networking Index: Forecast and Methodology, 2010-2015., [http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-481360.pdf](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf).
- [31] CIDR Report for 8 April 2013, <http://www.cidr-report.org/as6447/bgp-origins.html>.
- [32] Alexa Web Information Service, <http://aws.amazon.com/awis>.