The Aftermath of Prefix Deaggregation

Andra Lutu*†, Cristel Pelsser‡, Marcelo Bagnulo† and Kenjiro Cho‡
*Institute IMDEA Networks, Spain
†University Carlos III of Madrid, Spain
‡IIJ Innovation Institute, Japan

Abstract—Prefix deaggregation is recognized as a steady long-lived phenomenon at the interdomain level, despite its well-known negative effects for the community. The advertisement of more-specific prefixes provides network operators with a fine-grained method to control the interdomain ingress traffic. Moreover, customer networks combining this mechanism with selective advertisements may decrease their monthly transit traffic bill and potentially impact the business of their providers.

In this paper, we develop a methodology for Internet Service Providers (ISPs) to monitor new occurrences of prefix deaggregation within their customer base. Moreover, the ISPs can detect on their own when deaggregation may decrease the transit bill of their customer networks. We first examine the ISP's BGP routing data for new cases of prefix deaggregation generated by customers. Then, we check for selective advertisements of the newly generated prefixes using external routing data. We look beyond the incentives for deploying this type of strategy and instead we examine its economic impact. We exemplify the proposed methodology on a complete set of data including routing, traffic, topological and billing information provided by a major Japanese ISP and we discuss the implications of the obtained results.

I. INTRODUCTION

The Internet is an ever-growing and evolving ecosystem formed through the dynamic interconnection of many independently managed networks, also known as Autonomous Systems (ASes). The Border Gateway Protocol (BGP) is responsible for the exchange of reachability information and the selection of paths according to the routing policies specified by each network. The way in which the traffic flows between networks is influenced by the path dynamics triggered in the evolution of the Internet topology by failures, the maintenance of BGP sessions and, most importantly, the routing policies of each network [1], [2], [3]. By tweaking BGP configurations, ASes interact to coordinate the exchange of IP traffic according to various technical and economic necessities.

One important task achieved through the use of traffic engineering tools is the optimization of the routing function in order to allow the ASes to shift the traffic on its incoming links in the most effective way [1]. Address space fragmentation offers a high granularity for incoming traffic manipulation. This technique, also known as *prefix deaggregation*, allows networks to divide their assigned address blocks in different-sized sinks of traffic. Regardless of the strongly negative side-effects on the scalability of the global routing system [4], [5], [6], numerous operational reasons drive a part of the network operators to continue employing prefix deaggregation [6].

For example, in order to achieve load balancing over different incoming links or to attract traffic on cheaper links, ASes may combine prefix deaggregation with selective advertisements. By doing so, the corresponding incoming traffic for the more-specific prefixes only flows through the preferred transit provider towards the customer network. Also, by splitting their address space and advertising the smaller blocks to different regional providers, geographically-spread networks can divert different amounts of traffic corresponding to different points of presence (PoPs), thus attracting traffic into their network through the PoP closest to the final destination [1]. To assure the robustness of the routing system to failures, the deaggregating ASes usually inject, alongside the more-specifics, the less-specific covering prefixes to all providers. This type of behaviour is, however, costly for the provider. In addition to the need to support an increased number of prefixes in their routing tables, providers may also experience diminished revenues.

In this paper, we take the point of view of an ISP and ask a two-staged question: **(1) How extensive is the use of prefix deaggregation among the customer networks? (2) Does deaggregation combined with selective advertisements decrease the transit bill?**

The reasons for which networks deploy prefix deaggregation have been extensively addressed and documented by the research community [4], [5], [6]. Our study looks beyond the motivations for deploying such a technique and focuses on *detecting deaggregation within a predefined time-window and the additional possibility of economic impact* on the transit bill towards the providers.

Even if the majority of Internet routes is reportedly very stable in time [7], Teixeira et al. show in [8] that BGP routing changes are the main cause for the majority of the large variations in the traffic demands. Since the transit bill depends on the peak traffic usage and not on the total traffic usage, these variations translate in artificially larger monthly transit bills.

Previous work [9] analytically concludes that after combining the selective advertisements with the deaggregation technique, the AS can enjoy a reduction of the traffic fluctuations on the transit links. By inferring the routing changes from publicly available BGP routing data, Lutu et al. conclude that indirectly restricting the routing diversity towards a certain destination prefix through selective advertisements of more-specific prefixes translates into monetary savings for the deaggregating network. Consequently, the customer network

not only acts counter to the best recommended practices, but also may indirectly impact the business of its providers. This economic impact is enabled by the current operational status of the Internet, and in particular the widely used billing model based on the peak traffic usage [10], [11].

A. Contributions

Our main contribution in this paper is the methodology proposed to identify new cases of deaggregated prefixes by the customers of an operational ISP within a certain time-window. Recognizing as a reality the relatively stable usage of prefix deaggregation as a traffic engineering method in the Internet [6], we enable any operator with the necessary tools to detect which are the customers which are new deaggregators and monitor their behaviour in time.

We define *strategic*[1] *deaggregation* as the action of splitting the address block and selectively injecting each more-specific prefix to different disjoint subsets of providers. Customers which exhibit this behaviour may be able to game the 95^{th} percentile billing rule and possibly have a negative impact on the business of their ISPs.

We propose a passive measurement approach for the detection of strategic deaggregation events and assessing their economic consequences. The novelty of the approach is the manner in which it merges different types of information characteristic to an ISP in order to have a complete picture on the operations of its customer networks. This requires obtaining and processing routing, topology, traffic and billing information and molding it in order to reach the correct level of understanding on the impact different customers might have on their providers. Any ISP interested in detecting the occurrence of this phenomena within its customer base can build the dataset and apply the proposed methodology.

First, we use the ISP's BGP routing data to identify cases of prefix deaggregation among the customer networks. Second, we verify if any of the more-specific prefixes are selectively advertised by the customer networks. To this end, we check all the routing information gathered from ASes that are active monitors within the RIPE RIS [12] and RouteViews [13] projects. Third, in order to quantify the economic impact on the transit provider for the reported cases of strategic deaggregation, we also perform the traffic analysis. To approximate the decrease in the monthly transit bill, we require traffic data corresponding to the *strategic deaggregation* approach and to the situation prior to employing this method. On top of the traffic information separated for the before and after cases, we map the billing scheme used by the ISP.

Allowing the ISP to determine the deaggregation strategies deployed may raise awareness on the operations of its own customers. Even without the evaluation of the economic impact, this can be viewed as an alarm triggered for customers which may negatively impact the business of their provider.

---

[1]We use here the term *strategic* to accentuate the fact that the decision is based on optimizing behaviour, since it might increase the benefits for the network deploying it. This relies on definitions provided in rational choice theory.

The ISP can further monitor their evolution regarding the level of deaggregation deployed.

We exemplify the usage of the proposed methodology on a complete dataset provided by a major operational ISP. The dataset spanning a period of two months (May-June 2012) includes BGP routing information that enable us to monitor the dynamics of the customer address space over time. The number of customer networks advertising new more-specific prefixes within the period of analysis is relatively small: less than 10% of the customer networks were actively injecting new more-specifics within the analyzed period. After checking the external routing data, we detect one case of strategic deaggregation being deployed during the analyzed time-window. We perform the traffic data analysis for this case and we conclude that the customer AS may enjoy a reduction of 20% on its monthly transit bill.

The rest of the paper is structured as follows. in Section II, we present an example to intuitively explain how prefix deaggregation can impact the transit bill. In Section III, the toy example is properly formalized in the general Internet model. The methodology for detecting cases of the deaggregation is proposed in Section IV. In Section V we describe the dataset which we use for exemplifying the use of the proposed methodology and discuss the results obtained. In Section VI we present our conclusions and discuss future work.

II. TOY EXAMPLE

We introduce next a toy example to illustrate how a network changing its strategy from non-deaggregation to deaggregation can benefit from a decreased transit traffic bill, thus possibly impacting the revenues of its providers. For simplicity, let us consider the case of a network buying transit from two different providers, as in the scenario depicted in Figure 1.

We first analyze what happens with the traffic on the two different transit links before the customer deploys the deaggregation strategy. Let us assume that the destination network is advertising the same prefix $1.1.0.0/16$ over two different transit links, corresponding to the initial non-deaggregation behaviour represented in Figure 2.a. We limit the number of traffic sources at two, out of which one is generating $\frac{3}{4}$ of the whole traffic demand T (the *"three quarters"* source) and the other one, the rest (the *"one quarter"* source). Assuming the use of the 95^{th} percentile pricing model, in which the monthly bill is the function of the peak level of traffic, we monitor the level of traffic on each link during one month. We consider that the *"three quarters"* source is sending its traffic on link l_1 for more than 5% of the period, after which, due to a routing change, it starts forwarding its traffic on link l_2. The *"one quarter"* source suffers the opposite events, namely it switches from link l_2 to link l_1 for more than 5% of the billing period. As a result, because the traffic on each transit link has a level of $\frac{3T}{4}$ for more than 5% of the billing period, the chargeable amount of traffic for each provider is $\frac{3T}{4}$. Assuming c is the cost per unit of transit traffic, the total cost payed for the transited traffic T is $c\frac{3T}{2}$, which is $c\frac{T}{2}$ higher than the real cost $c * T$. We note that even if the traffic level on link l_2 has
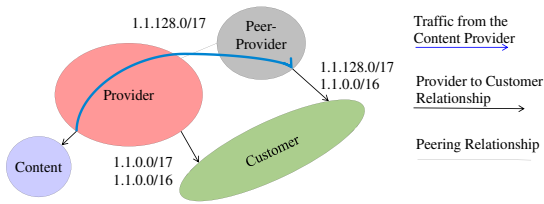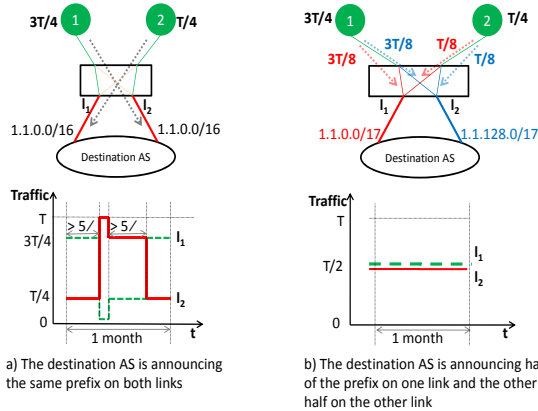
Fig. 1.   Strategic prefix deaggregation example.

a) The destination AS is announcing
the same prefix on both links

b) The destination AS is announcing half
of the prefix on one link and the other
half on the other link

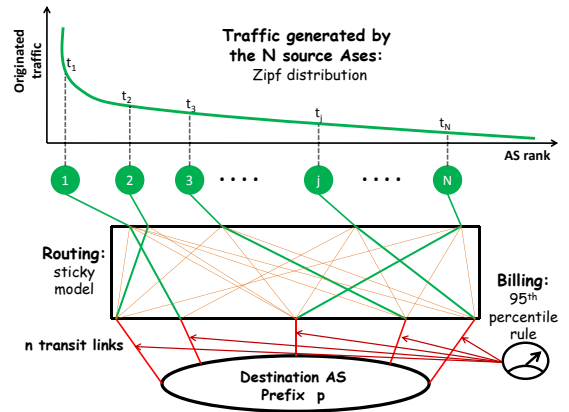Fig. 2.   Toy example representation.

Fig. 3.   Graphical representation of the proposed general Internet model. The sticky model for routing changes assumes that the sources stick to the green path from the initial-state with a probability of $1 - p$, and change to an alternative yellow path with a probability p.

III. OVERVIEW: GENERAL INTERNET MODEL

We overview here the settings of the general Internet model proposed for the study of the impact of prefix deaggregation on the transit traffic bill. By combining three important elements, i.e. the interdomain path changes, the 95^{th} percentile billing rule broadly used in today's Internet and the skewed distribution of traffic demand on the source networks, the model offers the underlying structure for the analysis of the phenomena associated with the deaggregating strategy.

Given that in the current Internet paths are calculated independently for each destination, we perform our analysis at the AS-level assuming, without loss of generality, the existence of one destination network and N sources of traffic. The model accounts for an AS with n symmetric transit links assumed to have the same capacity, which are accommodating traffic from N sources. For ease of presentation, a uniform distribution of incoming traffic on the destination address space is integrated in the model. In the case of an uneven traffic distribution, a correspondingly proportional prefix fragmentation can be found so that the amounts of traffic per more-specific prefix are comparable [9]. We first assume that the destination AS is advertising the same prefix on all its transit links.

The initial-state set of routes at the beginning of the billing time interval is achieved after first completing the selection process on the available BGP paths between each source AS towards the destination prefix. The selection process is modeled as random, but the model can be extended to accommodate any other type of process. The routing changes are modeled using a sticky process, where a probability p is assigned for the path changing between a source and the destination in terms of which transit provider is used relative to the initial-state.

The total traffic generated towards the destination is distributed among the N sources according to Zipf's law, as previously described in [14]. Given a ranking of the Internet entities, the Zipf law states that the traffic generated by a network is inversely proportional to its rank. This assumption is consistent with the traffic measurements in [15], as the Zipf distribution is a particular case of a power law distribution

a spike of T, it does not impact the transit bill, since it last less than 5% of the billing period.

If the destination AS deaggregates its prefix, it can avoid the fluctuations of traffic caused by routing changes and also the artificial augmentation in the transit traffic monthly bill. Consider, for example, that the destination AS in Figure 2 divides its address space into two more-specific prefixes and announces each on a separate link, as observed in Figure 2.b. This corresponds to deploying the strategic deaggregation behaviour. If we assume uniform distribution of incoming traffic for the prefix, each more-specific prefix now receives half of the traffic generated by each source. Consequently, the routing changes do not increase the 95^{th} percentile and the transit bill for the customer reflects the actual traffic demand.

The reasons for deploying the deaggregation strategy may include a wide variety. For example, when analyzing the topology from Figure 1, one reason might be the need to avoid the capacity upgrade on the link between the Customer and the Provider, or even the need of the Customer to receive traffic for the more-specific locally from the Peer-Provider, thus avoiding hauling the traffic within his own network. Studying the motivation for employing this mechanism is, however, out of the scope of our analysis. We focus instead on the impact of the deaggregation strategy on the provider. In the scenario from Figure 1, the Provider network may not only be loosing revenue from the decrease of the transit bill. There is also an additional cost implied by having to haul the customer traffic through its own network towards the Peer-Provider. Thus, by simply monitoring customers' behaviour and detecting the cases of strategic deaggregation, an ISP can identify the customers who may systematically impact in a negative way its revenues.
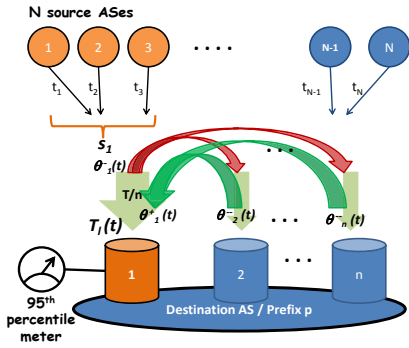
Fig. 4. Traffic dynamics for each transit link.

characterized by a skewness parameter $\alpha = 0.9$. One may argue that because of the large number of sources in the real Internet, small relative fluctuations of traffic can be expected. However, the skewness of the traffic distribution on sources has an important effect on the amount of traffic switching between transit links. This is justified by the fact that if a large source of traffic becomes instable, then it shifts important amounts of traffic between different routes, thus heavily impacting the traffic distribution on the incoming links towards the destination.

The traffic dynamics are captured in Figure 4, where we can observe how during the time of analysis, traffic may shift from one transit link to another. The routing changes imply a shift in traffic from one transit link to another, corresponding to the amount of traffic accommodated on the unstable source path. Consequently, the total volume of traffic on any transit link i, where $i = 1, ..N$, between the destination network and its N providers has the following expression, which changes at every time-slot t during the time of the analysis:

$$T_i(t) = \frac{T}{n} - \theta_i^-(t) + \theta_i^+(t), \qquad (1)$$

where $\theta_i^-(t)$ represents the traffic leaving link i and $\theta_i^+(t)$ represents the expected value of the traffic shifting from the rest of the links to link i. The expected value of the traffic is $\frac{T}{n}$, whilst the variation depends on the unstable amounts of traffic, i.e. the $\theta_i(t)$ traffic elements from (1).

The toy example suggests that these traffic fluctuation can be diminished by the use of prefix deaggregation combined with selective advertisements. In other words, the $\theta_i(t)$ traffic elements can be reduced to zero in the case where different parts of the prefix are injected to each of the transit providers. Consequently, the incoming traffic on each link is confined to the preferred incoming link through the selective advertisement. This further translates into monetary savings on the monthly transit bill, since the charging volume of traffic is implicitly decreased.

In the analysis on the operational Internet we focus on the two behaviours also presented in the toy example, since here we are bound to see the most important impact: 1) *no deaggregation*, where the address block is injected to all providers as the same prefix and 2) *strategic deaggregation*, where different parts of the prefix are injected to each of the transit providers. Nevertheless, intermediate deaggregation strategies

may be considered, where the resulting more-specific prefixes are injected to disjoint subsets of provider networks.

The billing method integrated in the model, i.e. the 95^{th} percentile billing rule, implies that the agreed billing period is sampled using a fixed-sized window, each interval yielding a value that denotes the traffic transferred during that period. A recent transit cost survey [16] has shown that the price per unit of transfered traffic, denoted here by c_t, decreases with the increase of the expected volume of transit traffic, following a concave dependency. However, this is only true when the increase of the expected amount of traffic is significant enough to justify the change. The authors show in [9] that the increase of expected traffic volume triggered by route changes does not require a change in the commit rate. Therefore the following linear cost function for the transit traffic is considered:

$$C = c_t * V \qquad (2)$$

, where V is the charging traffic volume (i.e. the 95^{th} percentile of the monthly traffic) of the destination AS and c_t is the corresponding transit traffic unit cost.

Consequently, the analytically derived formula for the relative monetary savings is a function of three parameters, reflecting the unique mixture between the three main elements: n - the number of links carrying chargeable volume of transit traffic, α - the skewness parameter of the distribution for traffic on sources and p - the path change probability included in the sticky model.

IV. THE METHODOLOGY

In this paper, we propose a novel methodology to identify new operational occurrences of more-specific prefix advertisements. This enables any ISP to monitor the amount of deaggregation generated by its customers and, in some cases, the impact it may have on its own revenues. The methodology is structured in three parts, each conveying relevant results concerning deaggregation dynamics within the customer base of an ISP. We summarize in Figure 5 the steps taken in the proposed methodology and show which type of information is required for each part.

Step 1: *Detect new more-specific prefixes*. First, we detect ASes which change their behaviour and start using deaggregation within a predefined time-window. For this step we require the BGP routing information from the ISP, as depicted in the first processing block depicted in Figure 5. We further expand on the mechanism in section IV-A.

Step 2: *Detect strategic deaggregation*. Second, we check for selective advertisements of the more-specific prefixes previously identified. As depicted in the second processing block from Figure 5, we use all the routing information from the monitors active in the the RIPE RIS and RouteViews projects.

Step 3: *Evaluate the economic impact*. Third, we try to determine if performing strategic deaggregation may lead to additional economic benefits for the customer network. For the cases of strategic deaggregation, we monitor the traffic data both (i) before deaggregation, when the address block is injected as one prefix to all providers (i.e. *no deaggregation*),
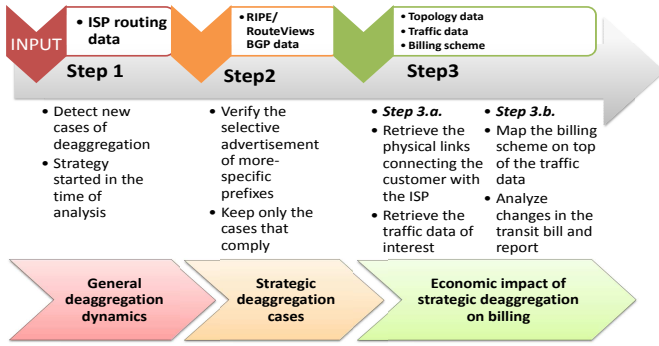
Fig. 5. The methodology steps: at each step we require a different input dataset depicted at the top of each processing block. At the bottom of each block, we can see the results we obtain at each step.

and (ii) after the strategic deaggregation, when the address block is fragmented into as many more-specific prefixes as the number of transit providers and each more-specific is selectively advertised to a different provider (i.e. *strategic deaggregation*). It is important to capture both these states, in order to be able to correctly quantify the economic impact of strategic deaggregation. We then evaluate the transit bill for each case and compare. This step is depicted in the third processing block from Figure 5.

Step 3.a: Extracting the traffic data on all the links connecting the provider with the identified customer which is deploying strategic deaggregation. This requires a previous mapping between customers and transit links from the ISP. We obtain the topology data from the router configuration files provided by the ISP. We observe this in the first sub-block of the third processing block in Figure 5.

Step 3.b: Finally, we move to *estimating the bill for the aggregated and deaggregated traffic patterns*. Thus, by applying the ISP's billing scheme to the traffic traces, we can quantify the impact of strategic deaggregation on the transit traffic bill. This step is depicted in the last processing block in Figure 5.

The methodology is aimed at working with a large and diverse collection of real data. Any ISP interested in detecting the occurrence of this phenomena within its customer base can build the dataset and employ the proposed methodology. Moreover, the tools we have developed can be made available for the research community upon request, as well as an anonymized version of the datasets used to exemplify the usage of the methodology.

A. Detection of Deaggregation Events

The detection algorithm we propose in *Step 1* for the identification of more-specific prefixes performs a comparative analysis of the BGP information obtained from the ISP. As depicted in Figure 6, the study on the dynamics of the prefixes advertised by the customer ASes during a particular time-window allows to separate the *newly injected* more-specifics prefixes, which first started to be advertised during the analyzed period.

We begin by choosing a *reference routing table*. The time-stamp of the reference routing table represents the *reference time*. The detection algorithm identifies the customer prefixes

based on the information from the provider (for example, customer routes are tagged with specific informational communities). We assume deaggregated prefixes exist at the reference time and look for the time deaggregation was first deployed. We progressively contrast the content of the reference routing table with each of the previous routing tables, starting from the furthest collection moment in time. This allows us to determine the presence of a covering prefix injected by the customer network and also the approximative moment of deaggregation. As depicted in Figure 6, we choose to verify the routing information from as much as one month before the reference time in order to capture the dynamics of prefix deaggregation in a timescale that is consistent with the billing period. However, the algorithm can be run on longer timescales (e.g. two months, three months, one year etc.), thus allowing the ISP to get a bigger picture on the deaggregation dynamics within its customer base at different timescales.

1) The Two-by-Two Routing Tables Comparison: We contrast the entries from the reference routing table with any other routing table collected in the period of analysis, to which we further refer as a *pair* routing table. We begin by first defining the set of prefixes present *only* in the reference routing table by separating the prefixes advertised only at the reference time and not present in the pair routing table, i.e

$$\Delta_i = P_{\text{ref}} - P_i \quad (3)$$

where P_{ref} represents the set of prefixes in the reference routing table and P_i, the set of prefixes installed in the paired routing table. For each of the prefixes in the Δ_i set defined above, we use a digital tree search [17] to identify the covering prefixes among the entries in the pair routing table. Assuming that no network is less specific than a /8, we are thus able to rapidly build the covering digital tree corresponding to each of the prefixes of interest. From each tree, we retrieve the least-specific prefix, i.e. the tree root, which we further use in the traffic data analysis. We do not examine the intermediate prefixes (shortly appearing intermediate phases in the deaggregation process), since for these there exists a more-specific prefix which can influence the manner in which traffic flows towards the destination.

By performing this comparative study using all the periodically collected routing tables from the ISP, we obtain an accurate picture of the evolution of the prefix deaggregation dynamics within the customer base of the provider. We monitor the changes of the previously defined prefix sets Δ_i during the analysis interval. The approximative time of deaggregation is, at the latest, the collection time of the first routing table snapshot which contains the candidate more-specific route known to already be installed in the reference routing table. This moment is marked in the time-line depicted in Figure 6 as the first moment where the more-specific prefix and the covering prefix are both present in the pair routing table.

B. Sifting the Results

In order to correctly identify the long-lived deaggregation events which may have an economic impact, we need to make
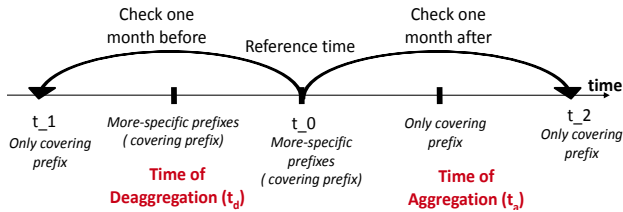
Fig. 6. Identifying deaggregated prefixes.

sure that the retrieved more-specifics are not sporadic events. For avoiding cases of dynamic deaggregation-aggregation behaviour, we filter out prefixes with intermittent presence in the routing tables, i.e. with a presence time lower than 5% of the billing period.

We apply the same detection algorithm to identify potential re-aggregation cases of more-specifics into their covering prefixes. We perform this latter step in order to assure that from the results provided by the algorithm we select only the more-specific prefixes that remain installed in the routing table for at least one month from the moment of deaggregation, and thus may impact the transit traffic bill. Past the reference time, the previously described two-by-two comparison algorithm actively detects cases of re-aggregated more-specific prefixes in the Δ_i set. We approximate the time of re-aggregation with the collection time of a pair routing table which contains only the covering prefix after the reference time.

1) Validation of Selective Advertisements: The selective advertisements validation process is further integrated in Step 2. We combine the internal routing view from the ISP with the external views taken from the ASes participating in the RIPE RIS and RouteViews project. In particular, we identify all the active providers used for reaching both the covering prefix and the more-specific prefix from Step 1.

We aim to check if the covering prefix is injected to all the active providers and the deaggregated prefix is selectively injected to only own. To this end, we analyze all the routing information retrieved during the corresponding time period (one month prior to the moment of deaggregation and one month after) from all the monitors whose routing tables we were able to retrieve from the public collectors. We monitor the routing information from each external AS towards the customer prefixes. Thus, we can infer the approximative number of active transit providers for the destination prefix by identifying the list of unique *second last-hops* (2LH) in the `AS-Path` BGP attribute after removing AS-Path prepending. The 2LH is the AS which we see before the destination AS in the `AS-Path`. This represents the provider used to reach the destination from the traffic source (i.e for some of the paths, this 2LH should be the ISP providing the data for this study).

We accept a certain error in the inferred connectivity degree of each customer, since we only have partial information on the interdomain routing. Given that the number of monitors active within the RIPE RIS and RouteViews project is limited, we have only a partial picture of how external sources of traffic reach the interest prefixes. However, since the sample of monitors is biased towards large Tier-1 networks, we assume that this is a reasonable approximation. We discuss how it

influences our results, along with other limitations of the methodology in Section V-C.

V. Exercising the Proposed Methodology

A. The Dataset

The measurements we perform in this paper use a complete dataset provided by a major Japanese ISP. The primary set of data we integrate in our study, the BGP routing data, is periodically collected from a monitor inside the ISP's network. Every two hours we obtain the complete routing information from the ISP. The routing snapshot (i.e. the complete BGP routing table taken at a certain moment in time) offers an accurate perspective on the dynamics of the customer prefixes which are of interest for our study. We assume that if a prefix is present in consequent snapshots it was also there between the snapshots. In addition, prefixes not present did not appear between the snapshots. The two-hours timescale offers a small enough granularity in order to capture the long-lived changes in the deaggregation strategy of the customer network. In order to correctly separate the *customer* network information from the BGP snapshots, we use the internal community tags the ISP uses for the routes received from its customers. We target only networks with public AS numbers, since it is likely that they also have multiple providers.

The collection of transit links through which each of these customers connects to the provider is necessary when extracting the traffic data corresponding to the detected cases of strategic deaggregation. In order to extract the topology information, we parse all the configuration files from the provider's edge routers, characteristic to different vendor-specific operating systems.

The traffic data is collected in NetFlow format and spans over a two months period i.e. May-June 2012. The sampling rate used for most routers is $\frac{1}{8,192}$. However, for some routers this may differ, depending on the traffic load on the router and its processing power. We analyze the traffic data that corresponds to the two different billing-compatible time intervals, i.e., one month before and another month after the deployment of strategic deaggregation. This limits us to detecting cases of customer networks deploying the deaggregation mechanism in the time-window corresponding to the two months of the study, i.e. May-June 2012. This limitation comes from the characteristics of the major ISP itself, which stores the traffic data for its customer only during the latest two months.

Finally, we add to our analysis the type of billing scheme employed by the ISP. Generally, the billing method relies on the 95^{th} percentile rule and the exact interval used for billing is the calendar month.

B. The Results

We illustrate the use of the proposed methodology on the dataset described in the previous section. First, we perform an extended analysis of deaggregation over a period of 6 months, from May until October 2012. This is aimed at providing a better understanding of the dynamics concerning deaggregation within the customer base of the Japanese ISP. To

Month	No. of customer ASes	No. of more-specifics
May 2012	7	154
June 2012	1	3
July 2012	2	3
August 2012	6	19
September 2012	5	42
October 2012	2	12

Fig. 7. Study case: result identified using the proposed methodology.

this end, we use the routing data provided by the operational ISP for a period of 7 months (May -November 2012). We iteratively select as a reference time every last snapshots taken each month from May to October 2012. By applying the algorithm described in Section IV-A, we are able to identify the set of customer networks that start to deploy deaggregation within the month previous to each of the reference times. We summarize the detection results in Table I. For example, we note that during August 2012 there were 6 different customer ASes which started to inject 19 new more-specific prefixes to the Japanese provider. We conclude that, generally, there are few customers deaggregating. And even more, the number of more-specifics injected to the ISP for each of the months analyzed is generally low, as observed in the third column from Table I. Overall, we observe 212 new more-specifics being injected throughout the 6 months analyzed.

Given that the traffic data is available only for May-June 2012, we present the analysis of the economic impact for deaggregation strategies identified in this particular period. In order to differentiate the cases of strategic deaggregation, we merge the results of the previous analysis with the external routing data from the monitors active in the RIPE RIS and RouteViews projects. We use the results corresponding to the prefixes deaggregated in May 2012, which also persist in the routing table for the next month.

Overall, we detect *154* more-specific prefixes[2] injected by the customers of the Japanese ISP during the month of May 2012. The prefixes are injected by 7 of the networks purchasing transit from the Japanese provider, as noted in Table I. Among the 154 more-specific prefixes first injected in May 2012, we are able to identify one case of deaggregation combined with selective advertisements, which fulfills all the requirements imposed. Our analysis shows that on the 28th of May, at around 16:00 hours, a customer prefix is deaggregated and the resulting more-specific prefix is injected to only one of the providers (i.e. the major ISP providing data). Moreover, the more-specific prefix is not re-aggregated into his covering prefix at any point during the following month of June 2012.

For the quantification of the impact of strategic deaggregation on the transit bill, we compare the traffic pattern for the identified prefix during a month prior to the moment of deaggregation (i.e. May 2012) with the traffic pattern

[2]The number of more-specifics injected in May 2012 is larger that in the other months due to a heavy deaggregator, which injects 120 more-specifics out of the total identified.
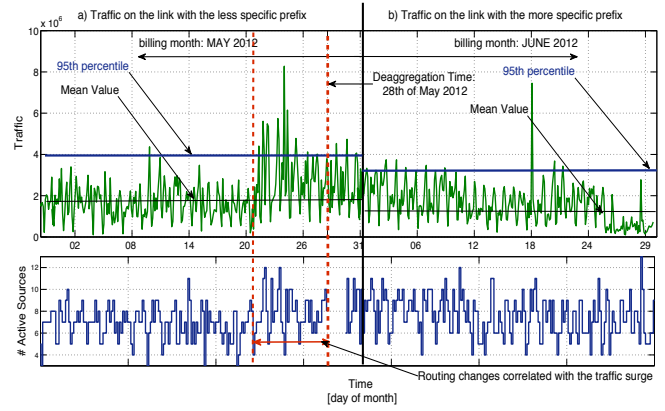
for the more-specific during a month after the moment of deaggregation (i.e. June 2012). Since the billing period used by the Japanese ISP is the exact calendar month, we compare the bill from May 2012 with the bill from June 2012. In order to extract from the traffic collection the data that interest us, we must first identify the physical links connecting the customer network under study and the provider. By parsing all the router configuration files, we obtain the identity of all the interfaces on the routers connecting the two networks. We then evaluate the chargeable amount of traffic for each case using the 95^{th} percentile billing rule. We conclude that, even if the expected amounts of traffic for the two prefixes are comparable, the transit bill is **20% lower** for the customer AS after selectively injecting the deaggregated prefix, as observed in Figure 7.

The difference in the chargeable volume of traffic per month may be due to the surge we observe in the traffic profile depicted in Figure 7 during May 2012. In order to check that this increase is caused by routing changes that influence the way large sources send their traffic towards the destination AS, we would need a complete view of the evolution in time of the BGP routing tables for the source networks. However, this type of information is unavailable at this point. Instead, we observe the changes in the number of active sources out of the top 20 which forward their traffic to the destination prefix via the Japanese provider, as depicted in Figure 7. We extract this information from the NetFlow traffic data of the Japanese ISP. The analyzed sources are prefixes with length 24 and are responsible for more than 50% of the total traffic towards the destination prefix. After the injection of the more-specific prefix, the traffic has a more stable behaviour than in the previous case and, also, the number of active traffic sources is more stable in time. We can also notice that there is a symmetry between the surge of traffic and an increase in the number of sources that forward their traffic through the transit link. The observed correlation between routing changes and traffic fluctuations supports the hypothesis according to which the 95^{th} percentile billing rule can be gamed by the customer networks by restricting the choices of transit links diversity towards the destination prefix. However, we cannot demonstrate the causality between the changes we observe

in the traffic pattern and the deaggregation strategy being deployed because of the lack of interest cases.

Based on the single perfect match for the strategic deaggregation strategy previously identified, we can only conclude that the study result supports the analytic observations for the economic impact of deaggregation. Thought the results presented in [9] may be true, the only case exposed using the proposed methodology is not sufficient for a generalization.

C. Limitations of the Methodology

Challenge 1: *Obtaining the data that enables a complete analysis of strategic deaggregation.* The quality of the results is conditioned by the quality of the data. Though the amount of information we handle is very large, it does not offer perfect information regarding the operations of the customers.

Challenge 2: *Dealing with partial information.* The Japanese ISP maintains fine-grained traffic information for its customer prefixes only for the latest two months from the moment of analysis. Since we require the traffic traces both before and after the strategic deaggregation mechanism was deployed, this limits the traffic analysis only to cases of strategic deaggregation that have occurred as far as one month previous to the moment of analysis.

Also, given that the number of monitors active within the RIPE RIS and RouteViews project is limited to approximatively 100, we have only a partial picture of how external sources of traffic reach the prefixes identified. Consequently, a prefix may be thought to be selectively advertised when it is in fact advertised to multiple providers. This happens because not all the paths appear at the available monitors. In this case, though, we should not see a lower transit bill than in the aggregated case.

Challenge 3: *Generalization the findings.* The results of our study show that the customers of the Japanese ISP do not make an extensive use of prefix deaggregation in general, and even less in the strategic form defined in this paper. The matching case exposed is not enough to draw any conclusions regarding the validity of the results presented in [9] and to generalize this type of behaviour. However, it is important to note that this may not be the case for every ISP.

VI. CONCLUSIONS AND FUTURE WORK

In this paper we propose a novel methodology for identifying cases of prefix deaggregation generated by the customers of any ISP within a predefined time-window. We focus the methodology to identify cases of selectively advertised deaggregated prefixes. We explain how the economic side-effect of the strategic deaggregation can be measured. In order to identify real occurrences of the interest phenomena, we demonstrate the use of this methodology on the real traffic and routing data from a major Japanese ISP. Overall, we do not observe much deaggregation generated from the customer networks of the ISP. We distinguish and analyze a strategic deaggregation case that fulfills all the constraints imposed in the methodology. We find that through selectively injecting more-specifics, the customer AS is able to smoothen the traffic

variations and save approximatively 20% on its transit bill. On a long term, this may negatively impact the business of the ISP. This result supports the hypothesis of an economic impact of strategic deaggregation, but it is not sufficient for generalizing it. For future work, we plan to expand our study to datasets spanning over longer periods of time. By doing so, we expect to add statistical relevance to the results and provide the validation for the economic impact of prefix deaggregation.

Until now, the transit providers did not have any incentives to refrain from advertising the deaggregated prefixes as injected by their customer [18], since the marginal cost is limited to an additional entry in an already bloated routing table. Provided the validity of the above-mentioned economic impact, the new incentives might be enough to push providers to change their strategy and transfer some of the costs of prefix deaggregation back to their customers. This could imply an important shift in the prefix deaggregation strategies adopted by the ASes in the Internet, moving the set of individual deaggregation strategies closer to the social welfare, where everybody enjoys increased benefits.

REFERENCES

[1] B. Quoitin, C. Pelsser, L. Swinnen, O. Bonaventure, and S. Uhlig, "Interdomain traffic engineering with bgp," *Communications Magazine, IEEE*, vol. 41, no. 5, pp. 122–128, May 2003.

[2] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja, and X. Xiao, "Overview and Principles of Internet Traffic Engineering." RFC 3272, 2002.

[3] N. Feamster, J. Borkenhagen, and J. Rexford, "Guidelines for interdomain traffic engineering," *SIGCOMM Comput. Commun. Rev.*, vol. 33, October 2003.

[4] G. Huston, "Analyzing the Internet BGP Routing Table," *The Internet Protocol Journal*, vol. 4, no. 1, 2001.

[5] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz, "BGP routing dynamics revisited," *SIGCOMM Comput. Commun. Rev.*, vol. 37, March 2007.

[6] L. Cittadini, W. Muahlbauer, S. Uhlig, R. Bush, P. Francois, and O. Maennel, "Evolution of Internet Address Space Deaggregation: Myths and Reality," *IEEE Journal on Selected Areas in Communications, special issue on Internet Routing Scalability*, 2010.

[7] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "Bgp routing stability of popular destinations," in *Proceedings of IMW'02*, 2002.

[8] R. Teixeira, N. Duffield, J. Rexford, and M. Roughan, "Traffic Matrix Reloaded: Impact of Routing Changes," in *PAM'05*, 2005.

[9] A. Lutu, M. Bagnulo, and R. Stanojevic, "An Economic Side-Effect for Prefix Deaggregation," in *NetEcon '12*, 2012.

[10] X. Dimitropoulos, P. Hurley, A. Kind, and M. P. Stoecklin, "On the 95-Percentile Billing Method," in *Proceedings of PAM'09*, 2009.

[11] R. Stanojevic, N. Laoutaris, and P. Rodriguez, "On economic heavy hitters: shapley value analysis of 95th-percentile pricing," in *Proceedings of IMC'10*, 2010.

[12] "RIPE RIS Raw data." [Online]. Available: http://www.ripe.net/data-tools/stats/ris/ris-raw-data

[13] "University of Oregon Route Views Project." [Online]. Available: http://www.routeviews.org/

[14] A. Dhamdhere and C. Dovrolis, "An agent-based model for the evolution of the internet ecosystem," in *Proceedings of COMSNETS'09*, 2009.

[15] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," *SIGCOMM Comput. Commun. Rev.*, vol. 40, August 2010.

[16] B. Norton, *The Internet Peering Playbook : Connecting to the Core of the Internet*. Dr. Peering Press, August 2011.

[17] D. E. Knuth, *The Art of Computer Programming, Volume 3: Sorting and Searching (2nd Edition)*. Addison-Wesley, 1998, vol. 3.

[18] C. Kalogiros, M. Bagnulo, and A. Kostopoulos, "Understanding incentives for prefix aggregation in BGP," in *Proceedings of the 2009 workshop on Re-architecting the internet*, ser. ReArch '09, 2009.