An Economic Side-Effect for Prefix Deaggregation

Andra Lutu
Institute IMDEA Networks, Spain
University Carlos III of Madrid, Spain

Marcelo Bagnulo
University Carlos III of Madrid, Spain

Rade Stanojevic
Institute IMDEA Networks, Spain
University Carlos III of Madrid, Spain

Abstract—**The injection of artificially fragmented prefixes through BGP is a widely used traffic engineering technique. In this paper we examine one particular economic side-effect of deaggregation, namely the impact on the transit traffic bill. We show that the use of more-specific prefixes has a traffic stabilization side-effect which translates into a decrease of the transit traffic bill. We propose an analytical model in order to quantify the impact of deaggregation on the transit costs. We validate our results by means of simulations and through the extensive analysis of real BGP routing information data.**

Index Terms—**BGP, Economics, Modeling, Traffic Engineering**

I. INTRODUCTION

The Internet is the interconnection of over 36000 domains known as Autonomous Systems (ASes), which engage in dynamic relationships that interplay with their technical and economic necessities. The routing between ASes relies on the Border Gateway Protocol (BGP), which is responsible for the exchange of reachability information and the selection of paths according with the policies specified by each domain.

The way in which the traffic flows in the interdomain is influenced by the path dynamics triggered in the continuous evolution of the Internet topology and the routing policies of each network. Hence, individual network managers need to permanently adapt to the interdomain routing changes and, by engineering the Internet traffic, optimize the use of their network. One important task achieved through the use of traffic engineering tools is the control and optimization of the routing function in order to allow the ASes to shift the incoming and outgoing traffic in the most effective way.

The injection of *more-specific prefixes* through BGP represents a powerful traffic engineering tool which offers a fine-grained method to control the interdomain ingress traffic. This technique implies that ASes selectively announce distinct fragments of their address block to different upstream providers. This type of phenomenon is commonly known as *prefix deaggreagation*. The most important negative side-effect of the widespread adoption of this technique is the artificial inflation of the BGP routing table [1], which can affect the scalability of the global routing system.

In this paper we study the impact of address-space fragmentation on the transit bill of the networks originating the more-specific prefixes. As a result of the unique interaction between the path changes in the current Internet [2], the distribution of traffic on sources and the widely used 95^{th} percentile billing method [3], [4], we find that the deaggreagating ASes enjoy one particular benefit from fragmenting their address space:

the decrease of the transit traffic bill. We propose an Internet model to analyze the cost for transit in two extreme cases of deaggreation, i.e. *no deaggregation* and *full deaggregation*[1]. We demonstrate that by using deaggreagation and scoped advertisements, the originating AS reduces the path diversity towards the injected prefixes. Thus, the amount of traffic destined to a particular sub-block of addresses is bound to the incoming link on which it was injected. This eliminates the possibility for any traffic fluctuations due to routing changes towards that particular prefix. Since the transit bill depends on the peak traffic usage and not on the total traffic usage, avoiding traffic fluctuations implies a lower monthly bill.

We begin, in section II, by explaining the intuition behind the network dynamics captured in the paper through the analysis of a toy example. We continue by presenting in section III the general model used for quantifying the effect of different deaggregating strategies on the ASes' transit costs. In section IV we estimate the model parameters by performing an extensive analysis of real BGP routing information. We afterwards quantify the actual impact of deaggregation. We further contrast the model-generated results with simulation and data-driven results. Finally, we conclude the paper in section V and present directions for future work.

II. TOY EXAMPLE

In order to provide the intuition behind the BGP phenomena modeled and analyzed in this paper, let us consider the simple case of one destination network announcing the same prefix $1.1.0.0/16$ over two different transit links, like we can see in Figure 1.a. For this toy example, we reduce the number of sources in the interdomain to only two, out of which one is generating $\frac{3}{4}$ of the whole traffic T consumed by the destination network, and the other one, the rest.

We monitor the traffic on each transit link during one month. We consider that the $\frac{3T}{4}$ source AS is sending its traffic on link l_1 for half of the period, after which, due to a routing change, it starts forwarding its traffic on link l_2. The $\frac{T}{4}$ source AS suffers the opposite events, namely it sends the traffic during the first half of the month through link l_2 and for the second half of the month it switches to link l_1. The transit traffic cost is calculated using the 95^{th} percentile rule for each of the two links. As a result, because the traffic on link l_1 had a level of $\frac{3T}{4}$ for more than 5% of the billing period, the transit traffic

---

[1]We define *full deaggregation* as splitting the address block into as many more-specific prefixes as the number of different transit providers.
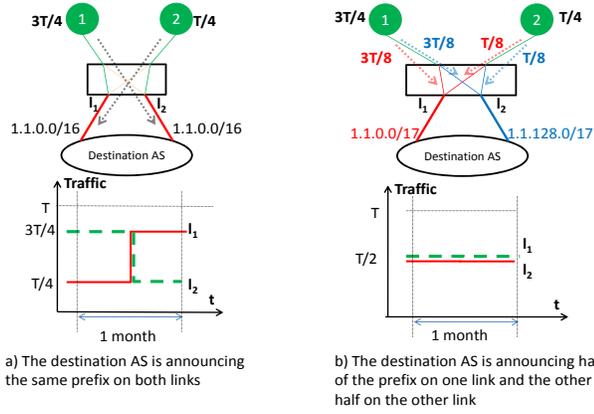
Fig. 1. Toy example representation.

Fig. 2. Graphical representation of the proposed Internet model for routing changes.

bill for link l_1 is $c\frac{3T}{4}$. Similarly, the transit traffic bill for link l_2 is also $c\frac{3T}{4}$, because for more than 5% of the billing period the traffic level was $\frac{3T}{4}$. Therefore, the total cost payed for the consumed traffic T is $c\frac{3T}{2}$, which is with $c\frac{T}{2}$ higher than the cost cT paid based on the 95^{th} percentile rule if no routing changes would happen.

We show in this paper that the destination AS can avoid the fluctuations of traffic due to routing changes through deaggregation and thus also avoid the implicit augmentation in the transit traffic monthly bill. Consider that the destination AS divides its address space into two more-specific prefixes and announces each on a separate link, i.e. announces $1.1.0.0/17$ through link l_1 and $1.1.128.0/17$ through link l_2, like we can observe in Figure 1.b. This means that no traffic fluctuations induced by routing dynamics exist on either links and, consequently, the sum of the 95^{th} percentiles on the two links is simply T. Consequently, in this scenario, the routing changes cannot impact the 95^{th} percentile and the transit traffic monthly bill for the destination AS is cT.

Regardless of the simplicity of the example presented here, the toy model does illustrate the basic intuition behind the observed phenomena, where routing changes may increase the transit bill for the destination AS, without increasing its total incoming traffic. In the following section, we propose the general Internet model which extends this example to capture more complex aspects of the Internet.

III. MODEL DESCRIPTION AND SAVINGS ANALYSIS

The general Internet model described here accounts for the route dynamics which are responsible for large traffic shifts in the interdomain, like previously observed in [2]. Given that, in the current Internet, paths are calculated independently for each destination, we model the Internet at the AS-level assuming, without any loss of generality, the existence of one destination AS and N sources of traffic. We perform our analysis on the case of the destination network with n links accommodating the incoming traffic from N sources. We assume a symmetric model where all the links have the same capacity.For the ease of the presentation, we assume an uniform distribution of incoming traffic on the destination address space. In the case of an uneven traffic distribution, it
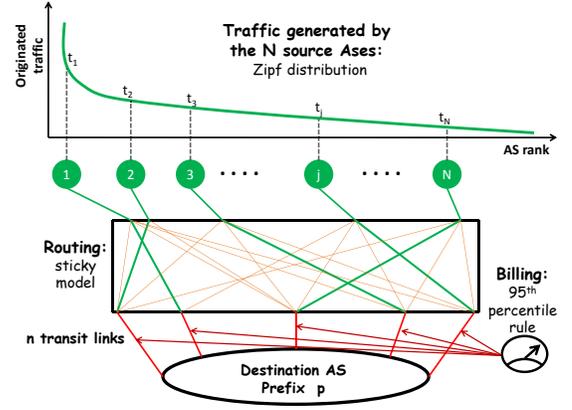
can be easily proved that a correspondingly proportional prefix fragmentation can be found so that the amount of traffic per more-specific prefix are comparable. As depicted in Figure 2, we integrate here three important elements, i.e. the interdomain path changes, the traffic model and the cost model. Their entanglement offers the necessary structure for studying the influence of deaggreagation techniques on interdomain traffic.

A. Deaggregation Strategies and The Sticky Model for Inter-domain Routing Changes

We model two extreme behaviours with respect to prefix deaggregation. We denote by λ the number of prefixes injected in the interdomain by the origin AS. One behaviour describes the AS which can choose to announce the same aggregated prefix through all its links ($\lambda = 1$). Alternatively, the origin AS can decide to evenly fragment its address block across its providers and announce disjoint more-specific prefixes on different transit links ($\lambda = n$). We assume that the assigned address space can be evenly split between the available number of links and announced by the destination AS as a single more-specific prefix separately on a different link[2]. Moreover, we assume that the announced prefixes are propagated as injected by the origin [5], which is aligned with current operational practices. Additionally, we assume the full reachability of the injected prefixes, meaning that every AS in the interdomain receives several routes for the n prefixes corresponding to the originating AS and selects one route for each prefix. In order to better understand the impact of the path changes on the cost for transit traffic, we analyze the path changes in a period of time relevant for the billing process, namely a month.

We model the interdomain route dynamics as follows. We model the initial-state set of routes as a random selection between the available BGP paths between each source AS towards any destination prefix, performed at the beginning of the analyzed time interval. In other words, we assume that if

---

[2]While it is not always true that the evenly divided address space can be announced only as a single aggregated prefix, we can find a particular fragmentation of the address space that would allow us to achieve the uniformity desired and announce the smallest number of more-specific prefixes.

a destination AS announces a prefix over n different links, then the probability that any of those links is a part of the forwarding path from a source towards the destination at the initial state is $\frac{1}{n}$. This assumption can be easily modified so that the model includes a different probability distribution on the available transit links.

In the rest of the time interval, we analyze the routing state by incorporating the dynamics of the routing process due to topology or policy changes[3]. We divide the month into 5 minutes intervals (consistent with the 95^{th} percentile billing rule) and use a time-slotted model to analyze the impact of path changes. Keeping in mind the fact that the majority of paths in the Internet are usually remarkably stable [6], we surmise in our model a small probability p for the path used in a given time-slot during a month to be different from the initially chosen one. Thus, we consider that the system is constantly performing a sticky process for readjusting to changes in the network.

B. Traffic model

In this section we model the traffic distribution on the available incoming links, depending on the manner the destination AS injects its prefix(es) in the interdomain.

We assume that each source network j included in our model generates an amount of traffic t_j towards a given destination in the interdomain, with the distribution depicted in Figure 2. We assume that the generated traffic t_j follows a Gaussian distribution characterized by the statistical mean μ_j and a variance σ_j^2.

1) Distribution of Traffic on Sources: We assume that the traffic generated towards a given destination is distributed among the existing sources according with Zipf's law, as previously described in [7]. This assumption is consistent with the traffic measurements in [8], as the Zipf distribution is a particular case of a power law distribution. Given a ranking of the Internet entities, the Zipf law states that the traffic generated by a network is inversely proportional with its rank. For any destination network we assign the following amount of incoming traffic from AS with rank j:

$$t_j = \frac{\frac{1}{j^\alpha}}{\sum_{k=1}^{N} \frac{1}{k^\alpha}} T = z_j(\alpha)T, \qquad (1)$$

where $z_j(\alpha)$ is the j ranked element corresponding to AS j in a Zipf distribution of N elements with the skewness parameter α. The total amount of traffic received by the destination AS can be expressed as the sum of all the traffic contributions $T = \sum_{j=1}^{N} t_j$, for all sources j in the Internet.

2) Distribution of Traffic on Transit Links: The total amount of traffic T consumed by a particular AS in the Internet consists of the contribution of all the sources in the interdomain. We analyze here the traffic distribution on the n ingress links of a destination AS.

[3]We do not consider the routing changes due to ingress link failures, as these changes cannot be accounted as potential savings since any operational viable deaggregation strategy must support backup links.
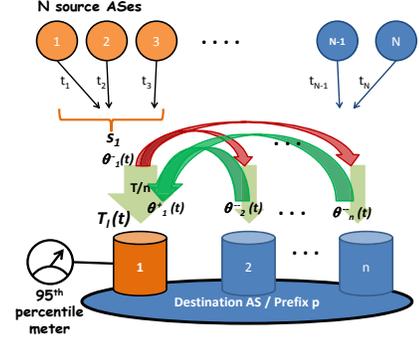
Fig. 3.   Traffic dynamics for each transit link.

We begin by characterizing the distribution of traffic on the incoming links of a destination AS that announces its address space as *one single aggregated prefix*. Consequently, any of the available links towards the destination network can be a part of the traffic forwarding path. For a given destination AS with n links we define the subset s_i of sources which have as initial state path a route which includes link i, where $i = 1, n$. For example, as depicted in Figure 3, subset s_1 includes all the source networks that have chosen transit link 1 in the initial phase of the model. Due to the fact that each link has the same probability of being chosen by each source for traffic forwarding in the initial state of the interdomain routing process, the expected value of the size of source sets s_i is of $\frac{N}{n}$ ASes. Consequently, in case of *no deaggregation*, the incoming traffic on each link in the initial state of the routing process has an expected value of $\frac{T}{n}$.

When dividing the month in many equal-sized time-slots, we further consider that the sticky routing model is adapting in each time-slot to the path changes which may have occurred. Therefore, in every time interval in the analyzed period, with a probability p the current forwarding path is different from the one used by the same AS in the initial state. This triggers the shift of a certain amount of traffic from link l to the rest of the links for the destination AS and the other way around.

We denote with $\theta_i^-(t)$ the random variable which represents *the traffic reduction* at moment t from link i and dividing among the rest of the transit links, as we can observe in Figure 3. The unstable traffic $\theta_i^-(t)$ leaving link i at moment t can be further expressed as $\sum_{j \in s_i} q_j(t)t_j$, where t_j represents the traffic generated by source AS j and has the expression in (1), s_i represents the set of sources with initial-state path including link i and q_j is either 1 if at moment t link i is a part of AS j's forwarding path or 0 in the contrary case. Formally, $P(q_j = 1) = p$ and $P(q_j = 0) = 1 - p$. The traffic reduction follows a Binomial distribution, i.e. $\theta_i^- \approx Binomial(\frac{N}{n}, p)$, with $i \in [1, n]$.

When analyzing the traffic on a link we also have to consider *the traffic increase* in the current link i from receiving traffic from the rest of the links $k \neq i$. This represents only a fraction of the total traffic moving away from any other link towards the current transit link. We denote with $\theta_k^-(t)$ the traffic leaving any link k, where $k \neq i$. Similar to the case of link i, we can express $\theta_k^-(t)$ as $\sum_{j \in s_k} q_j(t)t_j, k \neq i$, where q_j

is either 1 or 0 depending if at moment t link i is a part of the forwarding routed used by the source AS or not and t_j has the expression in (1). The traffic shift probability is equal to the probability of path change p, i.e. $P(q_j = 1) = p$. The total unstable traffic is represented by $\sum_{k \neq i} \theta_k^-(t)$. This amount evenly splits between all the $n-1$ equiprobable alternative links, including the analyzed link i. Consequently, the expected value of *the incoming traffic* denoted by $\theta_i^+(t)$ on link i is represented by the $\frac{1}{n-1}$ part of all the total unstable traffic, i.e. $\frac{1}{n-1} \sum_{k \neq i} \theta_k^-(t)$.

We can now express the total volume of traffic on each link towards the destination, which changes at every time-slot t like showed in the following expression:

$$T_i(t) = \frac{T}{n} - \theta_i^-(t) + \theta_i^+(t), \tag{2}$$

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.

Therefore, the expressions for the statistical mean and variance for the total traffic on link i when a single prefix is announced over all the available links are:

$$\sigma_i^2 = p(1-p)\left[\left(1 - \frac{1}{(n-1)^2}\right)\sum_{j \in s_i}^{|s_i| = \frac{N}{n}} t_j^2 + \frac{1}{(n-1)^2}\sum_{j=1}^{N} t_j^2\right];$$
$$\mu_i = \frac{T}{n}. \tag{3}$$

In the alternative strategy case of *full deaggregation*, the AS is announcing as many more-specific prefixes as number of transit links. Consequently, the size of the set of sources with a route that includes link i in the initial state is $|s_i| = N$. In other words, every source AS installs in its routing table a stable path for each transit link for the destination AS. This implies that the traffic shifting from one link to the others is zero and, similarly, the traffic incoming from the rest of the links is also null. Therefore, the variance of the traffic on each link resulting from route changes is $\sigma_i^2 = 0$, as the traffic forwarding paths are very stable. Consequently, the incoming traffic on each link equal with $\frac{T}{n}$ is confined to the preferred incoming link and does not fluctuate during the analyzed period.

C. The Cost Analysis

The 95^{th} *percentile rule* is currently the most widely-spread billing method among ISPs [3]. This method usually implies that the agreed billing period (usually a month) is sampled using a fixed-sized window, each interval yielding a value that denotes the traffic transferred during that period. The resulting intervals are sorted and the 95^{th} percentile of this distribution is used for billing [3].

A recent transit cost survey [9] 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 convex dependency. However, this is only true when the increase of the expected amount of traffic is significant i.e. one order of magnitude. In the case where the increase of expected traffic volume is in the same order as the initial traffic

volume, the cost per Mbps remains constant. We expect that the variations in traffic do not change the order of magnitude of the received traffic, therefore we consider the following linear cost function for the transit traffic: $C = c_t * V$, where V is the charging traffic volume (i.e. the 95^{th} percentile of the monthly traffic) of the destination AS i and c_t is the corresponding transit traffic unit cost. We consider that the total charging traffic volume for any destination AS represents the addition of all the chargeable traffic volumes on each incoming link, and therefore can be expressed as

$$V = \sum_{i=1}^{n} (\mu_i + 1.96\sigma_i), \tag{4}$$

where n represents the number of incoming links for the destination AS, and μ_i and σ_i have the expressions from (3). Given the fact that the traffic on link l follows a Binomial distribution $B(N, p)$, we can approximate it with a *Normal (Gaussian) distribution* $N(\mu_i, \sigma_i^2)$. The expression $\mu + 1.96\sigma$ from (4) represents the estimation of the 95^{th} percentile of a Normal random variable $N(\mu, \sigma^2)$ representing the individual traffic volume on the incoming links.

In order to capture the impact of deaggregation on the transit traffic bill, we evaluate the amount of chargeable traffic in the two previously mentioned cases. We calculate next the total amount of chargeable traffic on each link, i.e. the 95^{th} percentile of the link traffic, when no deaggregation is performed by the destination AS $(v_i|_{\lambda=1})$ and when the number of prefixes announced is equal to the number of available links $(v_i|_{\lambda=n})$:

$$v_i|_{\lambda=1} = \frac{T}{n} + 1.96\sqrt{p(1-p)}\sqrt{\sum_{j \in s_i} t_j^2 + \frac{1}{(n-1)^2}\sum_{k \neq i}\sum_{j \in s_k} t_j^2},$$
$$v_l|_{\lambda=n} = \frac{T}{n}. \tag{5}$$

Consequently, after substituting t_j with the expression in (1), the additional traffic on each link $\gamma_i = \gamma_i(p, \alpha, T) = v_i|_{\lambda=1} - v_i|_{\lambda=n}$ becomes $\gamma_i = 1.96T\sqrt{p(1-p)}\sqrt{\sum_{j \in s_i} z_j^2 + \frac{1}{(n-1)^2}\sum_{k \neq i}\sum_{j \in s_k} z_j^2}$ where $z_j = z_j(\alpha)$ represents the Zipf coefficient in (1). Furthermore, the difference in the total charging traffic volume for the analyzed destination AS with n links can be expressed as the sum of the traffic fluctuations in all the links, yielding the expression for the total volume of additional chargeable traffic:

$$\gamma = \gamma(n, p, \alpha, T) = \sum_{i=1}^{n} \gamma_i(p, \alpha, T). \tag{6}$$

The savings in transit traffic bill represent the cost c payed for the burstable, unstable traffic, i.e. $c = \gamma c_t$. Henceforth, the saved amount in the transit traffic bill represents a fraction of

$$RS = \frac{\gamma}{T + \gamma} \tag{7}$$

out of the actual price payed for the consumed traffic without deaggreagation. Substituting the expression for γ with the

expression in (6) yields that the relative transit traffic savings are a function of the number of links towards the destination AS, the instability probability p and the skewness parameter $\alpha : RS = f(n, p, \alpha)$.

IV. NUMERICAL RESULTS AND MODEL VALIDATION

In this section we first obtain numerical results for the analytical model. We then validate the model by contrasting these results with quantifications from both simulations and data driven estimations using real BGP traces.

In order to apply the proposed model to the current Internet and estimate the potential savings in the transit costs, we need first to assign realistic values to the model parameters, namely, N, n, α and p. Parameter N stands for the number of ASes in the Internet, which is in the order of 36000. The skewness parameter α for the Zipf distribution on the traffic sources is estimated in the current state of the art [7], [10] to a value of 0,9. As parameter p represents the probability of a change in the ingress link used by the source ASes to send traffic towards the destination, we estimate it by analyzing the data set containing real BGP routing information. Parameter n represents the number of transit links which we estimate in the next section following the routing data analysis.

A. Data Set

The data set used includes the full BGP routing table snapshots taken every 8 hours from66 different ASes present in the RIPE database [11], during the months of December 2010 until May 2011. This adds up to a total of more than 35000 snapshots of full routing tables, containing the BGP routing information from the 66 analyzed sources towards more than 350000 destination prefixes. We approximate the amount of traffic generated by each source by extracting from the Zipf distribution of traffic on the 36000 sources only the elements corresponding to the official ranks for the set of 66 different analyzed sources. We use the official CAIDA ranks assigned based on the data-set from January 2011 [12]. We estimate the number of different transit links per destination by identifying the unique *second last-hops*[4] (2LH) in the paths installed in the routing tables. We find that more than 93% of ASes have at most 7 transit providers.

B. Estimation of the Instability Probability

In order to estimate the transit link instability probability, we further observe the changes in the 2LHs of the AS paths towards the destination prefixes in the analyzed routing tables. For each source-destination AS pair, we calculate the probability that in a given interval the source AS is not using the link selected in the initial state towards the same prefix. For each of the 66 sources analyzed, we evaluate the relative time the source AS is not using the path announced in the first time slot of the analyzed period towards every destination prefix. Next, we match every destination prefix to the originating AS

[4]The *second last-hop* is the AS which we see before the destination AS in the `AS-Path` BGP attribute and it represents the upstream provider used by the source to reach the destination.
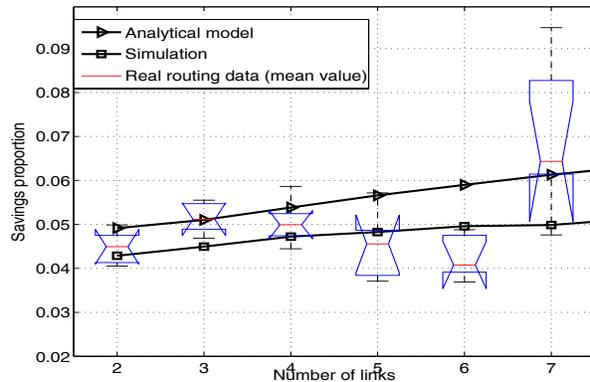
Fig. 4. Model generated, simulation and real-approximation savings curve for a deaggregating AS.

and average the time spent on an alternative path for all the prefixes announced by the destination AS. We thus obtain the probability that a source uses a path towards each destination AS which is different than the initial one. This implies that, for a proportion p of time, traffic may shift from the initial-state transit link towards another of the remaining equiprobable transit links. We approximate the parameter p with the mean value of the transit link instability probability over all the observed sources, yielding a value of $p = 3.5\%$.

C. Savings Quantifications using the Analytical Model

We observe in Figure 4 the model savings estimated for a destination AS depending on the number of links. We consider an instability probability of $p = 3.5\%$ in the interdomain and a Zipf distribution of traffic with 36000 elements and $\alpha = 0.9$. We observe that a destination AS with $n \in [2, 7]$ transit links may have an additional cost incurred by the route instabilities in the interdomain that can reach up to 6.5% of the transit traffic. According with the results presented in Figure 4, the average value of the savings for an AS with 2 transit links, representing 40% of all the ASes, is equal with 4.9% of the usual transit traffic bill.

D. Model Validation through Simulation

In order to check the accuracy of the results obtained analytically, we simulate the proposed model for a destination AS with $n \in [2, 7]$ transit links. We consider that the destination AS receives traffic from 36000 source ASes, following a Zipf distribution with skewness parameter $\alpha = 0.9$. We evenly divide the Internet traffic signal in 100 equal-sized time-slots for which we define the sample value of the traffic level, which is consistent with the number of routing snapshots we have during a month. We consider that in the first time-interval, each source AS uniformly chooses one of the n providers. In the remaining time-slots we simulate the sticky BGP selection algorithm by conceding a probability 3.5% of a different incoming link to be more appropriate for the source network than the initially chosen one . We define the 95^{th} percentile value for traffic on each transit link based on the values obtained through simulation. We apply next the formula in (7) to evaluate the savings due to the use of full deaggregation.

We represent in Figure 4 the curve of the average values of savings (estimated with less than 1% margin of error at 95% confidence level) for a destination with $n \in [2, 7]$ links. We thus observe an average value of 4.3% of the transit traffic bill savings for a destination AS with 2 different transit links.

The difference between the analytical model and the simulation comes from the fact that the analytical model uses for the 95^{th} percentile the Normal-based approximation confidence interval, i.e. $\mu + 1.96\sigma$, while this does not occur in the model simulation. In the latter case, we already have all the samples of the discrete Binomial Distribution, for which we can easily define the 95^{th} percentile of the link traffic level.

E. Savings Quantification using Real Routing Data

In this section we contrast the previously estimated numerical values for the transit link instability costs with approximations performed based on actual routing information. For this purpose, we process the BGP routing data present in the RIPE database corresponding to *6 months*, from December 2010 until May 2011 generated by 66 different sources. From comparing the routing tables we can evaluate the actual amount of routing changes towards a given destination AS in the Internet. Consequently, the path changes described by p are here substituted with genuine path changes inferred from comparing the routing tables over 6 months time. In our data analysis, we do not account for routing changes which are due to failures in the ingress links, since any operationally viable deaggregation strategy must support backup links. In order to filter out these cases, we remove from our analysis the destinations with a non-constant number of transit links present in each monthly data set. This approach is likely to remove a superset of the ingress-link failure cases, making our result to be only a *lower bound* of the potential savings.

When performing the reality based approximation of the savings, the Binomial approximated distribution of traffic is no longer needed, as we can infer the amounts of traffic on each link from evaluating the actual contribution of each source on every link. From the Zipf distribution with 36000 elements and $\alpha = 0.9$, we extract only the 66 elements corresponding to the sample of ASes.

We find that the amounts of savings are consistent over the 6 months, thus pointing out that the routing changes do impact the traffic levels in the same proportion each month. In Figure 4 we observe the savings estimated with real routing data for the destination ASes with $n \in [2, 7]$ transit links. The boxplot for each case shows the savings over the 6 months, where the central mark is the median, the edges of the box are the 25^{th} and 75^{th} percentiles, the whiskers extend to the most extreme data points not considered outliers. With 95% confidence level, the average amount of savings for an AS with 2 upstream providers lies in the confidence interval $[4.2\%, 4.77\%]$, which is consistent with the previous approximations.

V. Conclusions and Future Work

In this paper we have shown that injecting more-specific prefixes shows a particular monetary collateral benefit for the originating networks, namely a reduction of around 5% of their transit bill. Although the savings value might seem relatively small, it is in absolute terms non-negligible. Other traffic-engineering mechanisms (e.g. AS-Path prepending) could present with the same economic benefit, we focus here on deaggregation mainly due to the fine granularity with which it allows for traffic to be engineered. Several other phenomena adjacent to announcing more-specific prefixes have been identified and studied by the research community [13].

The result presented in this paper is the direct consequence of the current operational status of the Internet, as it relies on the unique mixture between the BGP-specific routing mechanism, the billing model and the difference in amount of generated traffic between different source networks. By varying any of the three parameters the model depends on, namely n (the number of transit links), α (the skewness parameter of the Zipf distribution of traffic on sources) and p (the path change probability), we can use the model to predict the evolution of the monetary savings in different scenarios.

For future work, we would like to measure the real economic side-effects of deaggregation for an actual operational network injecting more specifics in the current Internet. Thus, we aim to validate the analytical results with real-life values of the savings predicted.

VI. Acknowledgements

References

[1] T. Bu, L. Gao, and D. Towsley, "On characterizing BGP routing table growth," *Computer Networks*, vol. 45, no. 1, pp. 45–54, 2004.

[2] R. Teixeira, N. Duffield, J. Rexford, and M. Roughan, "Traffic Matrix Reloaded: Impact of Routing Changes," in *Passive and Active Network Measurement*, 2005.

[3] X. Dimitropoulos, P. Hurley, A. Kind, and M. P. Stoecklin, "On the 95-Percentile Billing Method," in *Proceedings of the 10th International Conference on Passive and Active Network Measurement*, 2009.

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

[5] 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.

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

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

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

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

[10] H. Chang, S. Jamin, Z. M. Mao, and W. Willinger, "An empirical approach to modeling inter-as traffic matrices," in *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, 2005.

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

[12] "Caida AS Ranking." [Online]. Available: http://as-rank.caida.org/

[13] P. Bangera and S. Gorinsky, "Impact of Prefix Hijacking on Payments of Providers," in *Proceedings of COMSNETS 2011*, January 2011.