

Volume-based Transit Pricing: Is 95 The Right Percentile?

Vamseedhar Reddyvari Raja¹, Amogh Dhamdhare², Alessandra Scicchitano³, Srinivas Shakkottai¹, kc claffy², and Simon Leinen³

¹ Texas A&M University, College Station, TX 77843, USA

² CAIDA/UCSD, San Diego, CA 92093, USA

³ SWITCH, Switzerland

Contact email: vamseedhar.reddyvaru@neo.tamu.edu

Abstract. The 95th percentile billing mechanism has been an industry *de facto* standard for transit providers for well over a decade. While the simplicity of the scheme makes it attractive as a billing mechanism, dramatic evolution in traffic patterns, associated interconnection practices and industry structure over the last two decades motivates an obvious question: is it still appropriate? In this paper, we evaluate the 95th percentile pricing mechanism from the perspective of transit providers, using a decade of traffic statistics from SWITCH (a large research/academic network), and more recent traffic statistics from 3 Internet Exchange Points (IXPs). We find that over time, heavy-inbound and heavy-hitter networks are able to achieve a lower 95th-to-average ratio than heavy-inbound and moderate-hitter networks, possibly due to their ability to better manage their traffic profile. The 95th percentile traffic volume also does not necessarily reflect the cost burden to the provider, motivating our exploration of an alternative metric that better captures the costs imposed on a network. We define the *provision ratio* for a customer, which captures its contribution to the provider’s peak load.

1 Introduction

The industry standard for transit billing is the 95th *percentile billing* method [1, 2] wherein a transit provider measures the utilization of a customer link in 5-minute bins over the duration of a month, and then computes the 95th percentile of these utilization values as the billing volume. The 95th percentile method has several attractive properties. First, this method is simple to implement, and uses data (e.g., SNMP) that the provider typically already collects. Second, it approximates the load that a customer causes the provider, while “forgiving” a few bursts (the top 5% of samples are ignored). While this transit billing method has remained fairly standard for over a decade, traffic patterns have evolved dramatically, from the dominance of client-server traffic in the early days of the Internet, to the rise and fall in popularity of peer-to-peer applications, to the rise of streaming video. Given that the traffic profile of a transit customer depends on the popularity of underlying applications, it is not clear that a transit billing scheme that may have been rational a decade ago is still appropriate.

In this work, we revisit the 95th percentile billing scheme from the perspective of a provider, to investigate whether this scheme approximately achieves its intended

objective of providing an easy-to-compute approximation of a customer’s traffic load to the provider. We first use 10 years of historical data from SWITCH, a Swiss research/academic network, and more recent data from 3 Internet Exchange Points (IXPs) to investigate how the 95th percentile of a customer’s traffic relates to: (1) its total traffic volume, (2) its nature as a predominantly inbound/outbound customer, and (3) its behavior as a heavy vs. moderate hitter. Second, we study the *fairness* of the 95th percentile scheme, and define a new metric called the *provision ratio* to investigate the relationship between the 95th percentile of customer and the contribution of that customer to the provider’s traffic load.

Analysis of these data sets reveals evidence that over the years the customers with a predominantly outbound traffic profile are able to maintain a lower 95th-to-average ratio than predominantly inbound customers, meaning that they have a lower billing volume for the same amount of traffic sent. Furthermore, the 95th-percentile pricing mechanism is *unfair*, because for many customers the 95th percentile may not reflect their cost burden to the provider, as there is little overlap between the customer’s peak and the overall (provider) peak traffic. Our results motivate the need to look for alternatives to the 95th percentile billing method that can better approximate a customer’s cost burden to the provider without adding too much additional measurement or computational overhead.

2 Datasets

SWITCH dataset: Our first dataset comes from SWITCH, a Swiss Research/Academic network which provides Internet connectivity to major universities and organizations in Switzerland. Currently, SWITCH connects about 50 research and education sites, acting as a transit provider for traffic that originates or is destined to those networks. SWITCH also provides connectivity to the public Internet via commercial providers, and hosts content caches of two large content providers. For traffic billing, SWITCH measures the utilization of each border router interface in both inbound and outbound directions in 5-minute intervals. To present a longitudinal analysis, we use historical datasets from SWITCH from January 2003 to December 2012.

IXP dataset: The second dataset consists of traffic statistics published by 3 Internet Exchange Points (IXPs) – Budapest Internet Exchange (BIX), Slovak Internet Exchange (SIX), and Interlan Internet Exchange (ILAN). These IXPs publish MRTG graphs with 5-minute utilization (inbound and outbound) for each network connected to the public peering fabric of the IXP. We collected these graphs every day for the month of August 2013 and used Optical Character Recognition tools [3] to parse them. BIX had 62 networks connected to its public peering fabric, while SIX and ILAN had 48 and 55 networks, respectively. Networks connect to IXPs to create (settlement-free) peering connections with other participating networks, and so the traffic statistics we see at an IXP are for a connected network’s peering traffic⁴. Castro et al. [3] showed that transit traffic and peering traffic have similar diurnal patterns and peak-to-valley ratios; in fact, the transit traffic for a network can be well-approximated as a multiplicative factor of the peering traffic. In our analysis we consider the IXP as proxy for a transit provider, and the networks connected to it as its customers.

⁴ While not explicitly disallowed, transit sale over the shared IXP fabric is rare [4]

3 Longitudinal Study of 95th percentile billing

We first describe two common methods of computing the 95th percentile traffic volume, and how the two methods can treat customers differently. We then classify networks based on two criteria: (i) major direction of traffic (inbound, outbound, and balanced); and (ii) volume of traffic (heavy-hitter and moderate-hitter), and present a longitudinal view of the traffic properties of these network types.

3.1 Calculation of 95th percentile

Although 95th percentile billing is the industry standard, there are two common implementations and several possible variations. The first method measures the inbound and outbound traffic in every 5 minutes over the month, calculates the 95th percentile for each direction, and uses the maximum of these two values. Most transit provider references to computing the 95th percentile use this method, e.g., [5, 6], so we use it in our subsequent analysis. The second method records the maximum of inbound and outbound traffic in each five minute interval, and calculates the 95th percentile value from the resulting data set. This second method seems to be less common although we found a few transit providers that bill using this method [7, 8]. The second method will yield a value greater than or equal to the first method, and the results will differ significantly for customers with balanced traffic profiles, but with inbound peaks occurring at different times from outbound peaks. We computed the 95th percentile for each network in the SWITCH dataset over 10 years. We found that the median ratio of the 95th percentile value for each network, computed using these two methods is close to 1, but the widest difference induces a 20% higher transit bill using the second method.

3.2 Classification of networks

Direction of Traffic: We divide networks into three categories based on the dominant direction of traffic. For each network, we measure the traffic that terminates within that network (inbound) and traffic that originates from that network (outbound). If the inbound traffic of the network is more than twice the outbound traffic we classify it as *heavy-inbound*, and if the outbound traffic is more than twice the inbound traffic we classify the network as *heavy-outbound*. Networks that do not satisfy either condition are classified as *balanced*. Typically, content providers are heavy-outbound, while eye-ball providers are heavy-inbound.

Volume of Traffic: We next classify networks based on the volume of traffic they generate/consume over a month into *heavy-hitter* and *moderate-hitter* networks. To define the two classes we evaluated the traffic contribution by the top 20% of networks in each month of the SWITCH and IXP datasets. The top 20% of networks consistently contributed between 80 and 90% of total traffic in the SWITCH dataset, and 75% of total traffic in the IXP dataset. Based on this observation, we classify the top 20% of networks in each month as *heavy-hitter networks* and the rest as *moderate-hitter networks*.

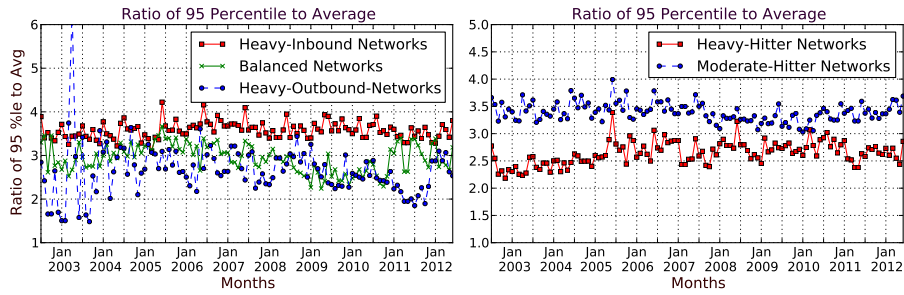


Fig. 1. Mean 95th percentile to average ratio for different network types in the SWITCH dataset. Heavy-inbound networks have a larger 95th percentile to average ratio than heavy-outbound networks. Also, moderate-hitter networks have a larger ratio than heavy-hitter networks.

3.3 95th percentile to average ratio

For each customer network, we first evaluate the 95th percentile to average traffic ratio; the average reflects the total volume of traffic, whereas the 95th percentile value gives an idea of the peak, and is also the traffic volume for which the customer is billed. If the two significantly differ, it suggests that the customer is paying primarily for its burstiness. Figure 1 shows the mean of the 95th percentile to average traffic ratio over time for networks in the SWITCH dataset classified by traffic direction and traffic volume.

First, we observe that the 95th percentile to average ratio has been fairly stable over the years for each type of network, despite the dramatic changes in overall inter-domain traffic patterns that have occurred during the same time. In the last 4 years, the mean ratio for heavy-outbound networks is between 2 and 3, while the mean for heavy-inbound networks is between 3.25 and 4. For balanced networks, the ratio is less than 3.25. Hence, heavy-inbound networks in general have higher 95th percentile traffic compared to heavy-outbound or balanced networks for the same average traffic. Consequently, heavy-inbound networks have a higher billing volume than heavy-outbound networks for the same amount of total traffic sent. We observe that the mean ratio is between 2.25 and 3 for heavy-hitter networks, especially in the last 4 years. However, the mean ratio always exceeds 3 for moderate-hitter networks in those 4 years.

Table 3.3 shows the mean 95th percentile to average ratio for different classes of networks in the IXP dataset. We observe that the mean ratio is higher for heavy-inbound networks than for heavy-outbound networks, consistent with our analysis of the SWITCH dataset. With the exception of BIX, the mean 95th percentile to average ratio for networks at the other two IXPs is larger for moderate-hitter networks than for heavy-hitter networks, meaning that moderate-hitter networks have a burstier traffic profile than heavy-hitter networks.

3.4 Skewness of the traffic distribution

The above analysis shows that heavy-inbound and moderate-hitter networks have a higher 95th-to-average ratio as compared to other networks, meaning that their traffic profile is likely to be burstier. Figure 2 illustrates the difference by plotting the mean skewness of the traffic distribution for each network type.

IXP	Heavy-inbound	Balanced	Heavy-outbound	Heavy-hitter	Moderate-hitter
SIX	2.6	-	1.7	1.4	1.9
BIX	2.82	2.3	2.1	2.59	1.94
ILAN	2.62	1.9	2.21	1.7	2.386

Table 1. Mean 95th percentile to average ratio for IXPs, using different network classifications. Heavy-inbound and moderate-hitter networks (except at BIX) generally have higher ratios.

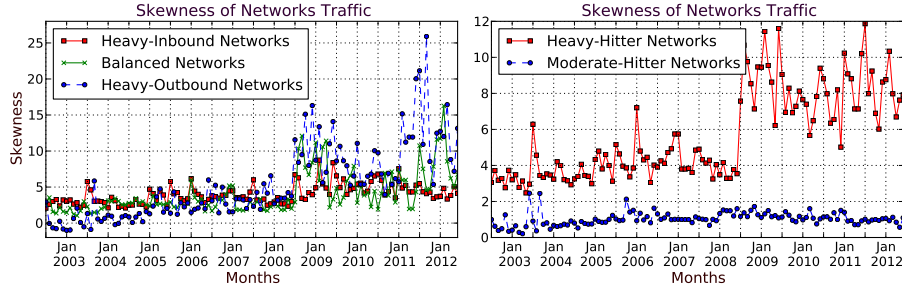


Fig. 2. Mean skewness for different network types in the SWITCH dataset. Heavy-outbound networks have a higher skewness, especially in the last 4 years. Heavy-hitter networks have larger skewness than moderate-hitter networks.

Skewness reveals how much the traffic distribution leans to one side of the mean; for a random variable X : $\text{Skewness} = E[(X - \mu)^3] / (E[(X - \mu)^2])^{3/2}$, where μ is the mean. If a probability distribution function is unimodal, then higher positive skew implies few values higher than the mean, i.e., the 95th percentile value would be closer to the average. The empirical probability mass function for the traffic of each network is unimodal for our data sets. Heavy-outbound networks have high positive skew (the mean is between 5 and 25), especially in the last 4 years⁵, compared to heavy-inbound networks or balanced networks, whose mean skewness is between 0 and 12 and 5 and 15, respectively. Similarly, heavy-hitter networks have higher positive skew than moderate-hitter networks. Table 3.4 shows the mean skew of traffic for networks at each IXP, classified according to dominant traffic direction and traffic volume. As in the SWITCH dataset, heavy-outbound and heavy-hitter networks generally have a larger skewness than heavy-inbound and moderate-hitter networks.

In summary, the 95th-to-average ratio has been stable for various classes of networks in our dataset over the last decade, indicating that a high-percentile billing scheme is still useful. Certain networks (particularly heavy-outbound and heavy-hitter networks) are able to achieve a lower 95th percentile to average ratio (perhaps using intelligent means of traffic shaping), and hence a lower billing volume for the same total amount of transit traffic. Traffic smoothing may allow networks to achieve a lower transit bill, but this says little about the contribution of those networks to the provider’s peak traffic. The 95th percentile of a network does not account for *when* the peaks occur, and so it is unclear whether it is fair to charge each customer using the same percentile.

⁵ The level shifts around 2009 coincide with SWITCH connecting to AMS-IX, acquiring hundreds of new peers, though the set of customers over which we compute statistics is unchanged.

IXP	Heavy-inbound	Balanced	Heavy-outbound	Heavy-hitter	Moderate-hitter
SIX	-0.56	-	0.04	0.3	-0.88
BIX	-1.6	-0.4	-0.19	-0.88	0.317
ILAN	-0.122	0.07	0.29	0.253	-0.11

Table 2. Mean skewness for networks in the IXP dataset. Heavy-hitter networks and heavy-outbound networks generally have higher skewness.

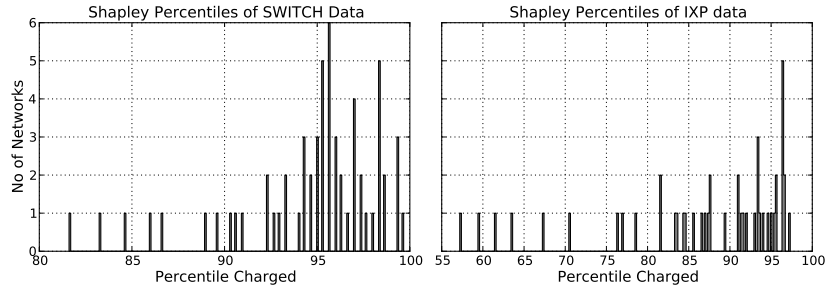


Fig. 3. Shapley value percentiles: SWITCH dataset (Mar 2012) and IXP dataset (SIX, Aug 2013).

4 Fairness of 95th percentile Billing

Motivated by the preceding discussion, we now focus on the fairness of the 95th percentile billing mechanism. We consider a billing mechanism fair if the amount of resources used by a network is reflected in the amount it is charged. An appealing idea in this context is the Shapley value, which assigns costs to the members in a cooperative game [9]. It possesses many attractive properties – it is *efficient*, i.e., the sum of costs assigned to each member is the total cost to the system, and it is *symmetric*, i.e., two members that have the same contribution will be assigned the same cost.

4.1 Shapley Value Percentile Billing

Stanojevic et al. [10] presented a model of the ISP cost allocation problem as a cooperative game. The cost function of a group is the 95th percentile of the total traffic obtained by adding the traffic of all members in that group. This cost estimate is consistent with the idea that the transit provider must provision for peak traffic, and is itself billed by its provider based on this value. The Shapley value (ϕ_i) of network i is then uniquely defined by $\phi_i = \frac{1}{N!} \sum_{\pi \in \Pi} (\mathcal{V}(S(\pi, i)) - \mathcal{V}(S(\pi, i) \setminus i))$ where \mathcal{V} is the cost function, Π is the set of all possible permutations of players \mathcal{N} and $S(\pi, i)$ is the set of all players in ordering π before i and including i .

Once we determine the Shapley value of each network, we need to map it to a billing percentile. Let the volume corresponding to the 95th percentile value of the total traffic be \mathcal{V} . Then (by efficiency) the Shapley values of the customer networks will satisfy $\mathcal{V} = \sum_i \phi_i$. Let the volume corresponding to the 95th percentile of network i be x_i .

Then the total volume billed by the transit provider under the 95th percentile billing scheme is $\sum_i x_i$, which we define as \mathcal{X} . Trivially, $\mathcal{X} \geq \mathcal{V}$. For an apples-to-apples comparison between the two billing schemes, we define the normalized Shapley value of network i as $s_i = \phi_i \mathcal{X} / \mathcal{V}$, so that the total billing volume in both cases is \mathcal{X} . Then each network can be charged based on a percentile that yields the traffic volume closest its normalized Shapley value, which is the “Shapley value percentile” of that network.

Computation of the Shapley value is quite complex—with N users, it has complexity order of $\mathcal{O}(N!)$. Even for a moderate size ISP, which has around 50 users, the complexity is of the order of 10^{64} . Stanojevic et al. [10] used a Monte Carlo approximation, which achieves a good trade-off between accuracy and complexity. We used this approximation to find the Shapley value percentile for the SWITCH dataset (month of March 2012) and the SIX IXP (August 2013). The results are shown in Figure 3. Clearly, the Shapley value percentiles are widely different from the 95th percentile .

In addition to computational complexity, the Shapley value percentile can be anywhere between 0 and 100. This approach lacks the ability of restricting the charging percentiles to a fixed range. The handicaps of directly using the Shapley value motivate a need for a simple proxy that captures its essence. A key observation is that a traffic profile has greater Shapley value when it is concentrated during the peak periods when demand is highest. Thus, Shapley value percentile billing would charge users with high peak traffic higher than users with off peak traffic.

4.2 Overlap rank

Building on the intuition developed in the last section that it is fair to charge more to networks with traffic during peak periods than off-peak periods, we will show how the current 95th percentile billing mechanism can lead to unfairness as it does not consider peak and off-peak periods. We define the peak periods of a transit provider as those in which the total traffic carried by the transit provider exceeds the 95th percentile of the provider’s total traffic. We similarly define the peak slots for customer networks. Based on the number of peak slots of networks that overlap with peak slots of the total traffic, we rank the networks from highest to lowest and call it the *overlap rank*. Thus, a network with rank 0 has the maximum number of peak slots that occur during the same time intervals as the peak slots of the transit provider. We also rank networks based on their 95th percentile and call it the 95th percentile rank.

Figure 4 plots overlap rank vs. percentile rank (normalized to 100) for the IXP dataset (first 3 plots) and one month (January 2012) from the SWITCH dataset (far right). If networks with high 95th percentile rank also had high overlap rank, most points would appear on the diagonal, and imply that 95th percentile billing is charging the contributors who necessitate the provisioning of large transit links. Figure 4 tells a different story. The points below the diagonal, especially those in the red shaded area (16% of networks for SWITCH) have a high 95th percentile rank but a low overlap rank, which means that their peaks are mostly in the peak period, but their billing volume is relatively lower. Analogously, the points above the diagonal line, especially in the gray region (15% of networks for SWITCH) correspond to low 95th percentile rank and high overlap rank. Their contribution to the peak period is low but they have a relatively high billing volume. Similar observations can also be made from the IXP graphs in Figure 4.

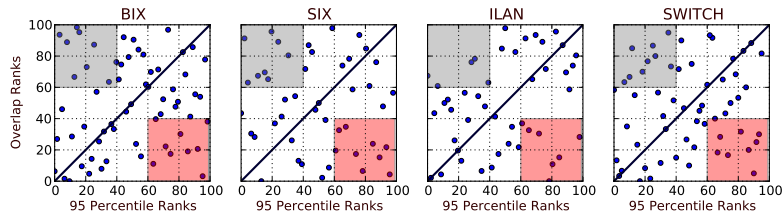


Fig. 4. Overlap rank vs 95th percentile rank for IXP dataset (Aug 2013) and one month of SWITCH dataset (Jan 2012). A large fraction of networks lie far from the diagonal, meaning they have a large billing volume but little overlap with the provider’s peaks, or vice versa.

4.3 Provision ratio

The overlap rank considers only the cardinality of overlap slots, without accounting for diverse traffic volumes. A good proxy for the Shapley value should capture the volume during peak slots, appropriately normalized with the amount of traffic generated by the network. We define the **provision ratio (PR)** of a network as the ratio of the average traffic during the peak slots of total traffic to the 95th percentile of that network’s traffic.

$$\text{PR of network } i = \frac{\text{Total traffic of network } i \text{ during peak slots} / \# \text{ of peak slots}}{95^{\text{th}} \text{ percentile of network } i\text{'s traffic}}.$$

The PR is essentially the ratio of traffic contributed by the network during the peak time slots (or average capacity provided to that network during these peaks) to the peak traffic of that network (excluding the top 5% of bursts); It can be viewed as the fraction of a network’s peak traffic that occurs during the provider’s peak periods. We propose that the PR can be an important component of a billing mechanism, because it captures the contribution of a network’s traffic to the provider’s peak. The PR is also robust to the exact thresholds used to compute it – we found that in our datasets, the provision ratio is robust to the exact threshold for defining a peak slot, e.g., if we change the 95th percentile to 85th percentile, the provision ratio does not change significantly.

The provision ratio is not equal to the Shapley value percentile in an absolute sense, but in a relative sense it appears to have the right characteristics. To quantify the similarity between the two, we find the percentage of orders preserved between all possible pairs of networks in both datasets. A transit provider with N customers will have ${}^N C_2$ customer pairs. For each pair, order is preserved if the network that is charged a higher Shapley percentile also has a higher provision ratio. We find that for the SWITCH dataset, the provision ratio preserves between 76% and 82% of orders in the SWITCH dataset (each month of 2012) and 89%, 75%, and 82% for the SIX, BIX, and ILAN IXPs, respectively (August 2013). The strong similarity of orders indicates that provision ratio is indeed order preserving.

4.4 Towards a new billing mechanism

One could argue that the 95th percentile billing scheme is an approximation, aiming for simplicity and predictability over fairness. At the other extreme is Shapley value pricing, which charges each user differently based on their actual contribution to the provider’s costs. An open challenge is how to achieve both objectives – fairness and

low computational complexity. We are currently exploring the use of the provision ratio in a scheme that determines the optimal percentile to charge a given customer. The objective of this scheme would be to vary the billing percentile per customer, and to use the provision ratio as a measure of the contribution of a customer to the provider's peak traffic. This pricing scheme would automatically assign lower billing percentiles (i.e., give discounts) to customers whose peak traffic does not contribute significantly to the provider's peak, and higher percentiles to customers that contribute most to the provider's peak. An important criterion for such a scheme is that the provider should be able to communicate information about its peak and off-peak periods to customers, without having to make its traffic profile available publicly. For this purpose, the provider could design a tool that accepts a customer's traffic profile and analyzes it in relation to its own traffic to determine the percentile at which it would charge the customer. Such a scheme would retain the attractive properties of burstable billing (because it is still based on a billing percentile), while better accounting for a network's contribution to total provider costs. Our initial investigation indicates that this problem can be formulated as a convex optimization, and hence solved efficiently.

5 Related Work

While network service pricing has been studied extensively, relatively little work has focused on specific mechanisms in the transit business, i.e., volume based pricing based on the 95th percentile rule. As early as 1999, Brownlee et al. [11] experimented with an alternative to the 95th percentile pricing mechanism, the "third quartile day", which they showed was a better estimate of the bandwidth requirements for customers of New Zealand's Kawaihiko network. Norton discussed 95th percentile pricing in his white papers, particularly the possibility of ISPs gaming the scheme to get free transit [12], and the impact of streaming video on the statistics of customer traffic [13]. Dimitropoulos et al. [2] studied the 95th percentile billing method using traffic traces, and investigated how the 95th percentile computed for a given network depends on factors such as the averaging window size and the effect of flow aggregation. In the context of broadband users, Stanojevic et al. [10] used the Shapley value approach to quantify the contribution of each broadband user to the total costs of the access provider. Valancius et al. [14] proposed that transit providers implement tiered pricing using just a few tiers based on the volume of traffic and the cost of carrying it to maximize their profits. However, their approach was targeted at properly structuring pricing tiers, i.e., the price per unit of traffic that the provider charges to a customer. The focus of our work is on the underlying traffic percentile at which a provider charges its customers.

6 Conclusions

In this paper, our goal was to empirically examine the effectiveness of the 95th percentile pricing scheme, using a decade of historical traffic data from a transit provider network and more recent data from three European IXPs. Our analysis shows that over the years, certain networks have lower 95th-to-average ratio than others – for the

datasets we studied, networks with predominantly inbound traffic have higher 95th-to-average ratios, and would incur a higher billing volume than those with predominantly outbound traffic (for the same amount of total traffic), and similarly for moderate hitters vs. heavy hitters. Furthermore, we find that the 95th percentile pricing scheme can be unfair, as the 95th percentile traffic of a network is often unrelated to the amount of time that network's peak traffic overlaps that of its provider, nor does it accurately represent the contribution of that network to the provider's peak traffic. We define a new metric, the Provision Ratio (PR) for a network, which is easy to compute and is able to capture the contribution of a customer traffic to the provider's peak.

Acknowledgements

We thank our shepherd, Sergey Gorinsky, and the anonymous reviewers for their constructive comments. This material is based upon work supported in part by NSF grants CNS-1149458, CNS-1017064 and a Cisco URP grant. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NSF or Cisco.

References

- [1] Odlyzko, A.: Internet pricing and the history of communications. *COMPUTER NETWORKS* **36** (2001)
- [2] Dimitropoulos, X., Hurley, P., Kind, A., Stoecklin, M.P.: On the 95-percentile billing method. In: Proceedings of PAM. (2009)
- [3] Castro, I., Stanojevic, R., Gorinsky, S.: Using Tuangou to reduce IP transit costs. *IEEE/ACM Transactions on Networking* (2013)
- [4] Norton, W.B.: Transit Traffic at Internet Exchange Points? Drpeering.net blog: <http://drpeering.net/AskDrPeering/blog>
- [5] G4 Communications: 95th Percentile Usage Billing Policy http://www.g4communications.com/docs/G4_95th_Percentile_Usage.pdf.
- [6] Axis Internet: 95th Percentile <http://www.axint.net/95th>.
- [7] AboveNet: Monitoring 95-percentile <http://john.de-graaff.net/wiki/doku.php/links/95-percentile>.
- [8] Cline Communications: 95th Percentile billing <http://clinecommunications.net/?ID\=33>.
- [9] Roth, A.E.: The Shapley Value: Essays in Honor of Lloyd S. Shapley. (1988)
- [10] Stanojevic, R., Laoutaris, N., Rodriguez, P.: On Economic Heavy Hitters: Shapley Value Analysis of 95th-percentile Pricing. In: Proceedings of IMC. (2010)
- [11] Brownlee, N., Fulton, R.: Kawaihiko and the Third-Quartile Day. *IEEE Communications* **38**(8) (2000)
- [12] Norton, W.B.: Transit Tactic - Gaming the 95th Percentile. Drpeering.net white paper: <http://drpeering.net/>
- [13] Norton, W.B.: Video Internet: The Next Wave of Massive Disruption to the U.S. Peering Ecosystem. Drpeering.net white paper: <http://drpeering.net/>
- [14] Valancius, V., Lumezanu, C., Feamster, N., Johari, R., Vazirani, V.V.: How many tiers?: pricing in the Internet transit market. In: Proceedings of ACM SIGCOMM. (2011)