Volume Based Transit Pricing: Is 95 the right Percentile?

Vamseedhar Reddyvari[1] Amogh Dhamdhere[2] Srinivas Shakkottai[1]
K Claffy[2] Alessandra Scicchitano[3] Simon Leinen[3]

[1]Texas A&M University

[2]CAIDA, UCSD [3]SWITCH

March 9, 2014

95^{th} Percentile billing

- 95^{th} Percentile billing has been the industry standard for charging customer networks of ISPs for well over a decade
- 95^{th} Percentile calculation

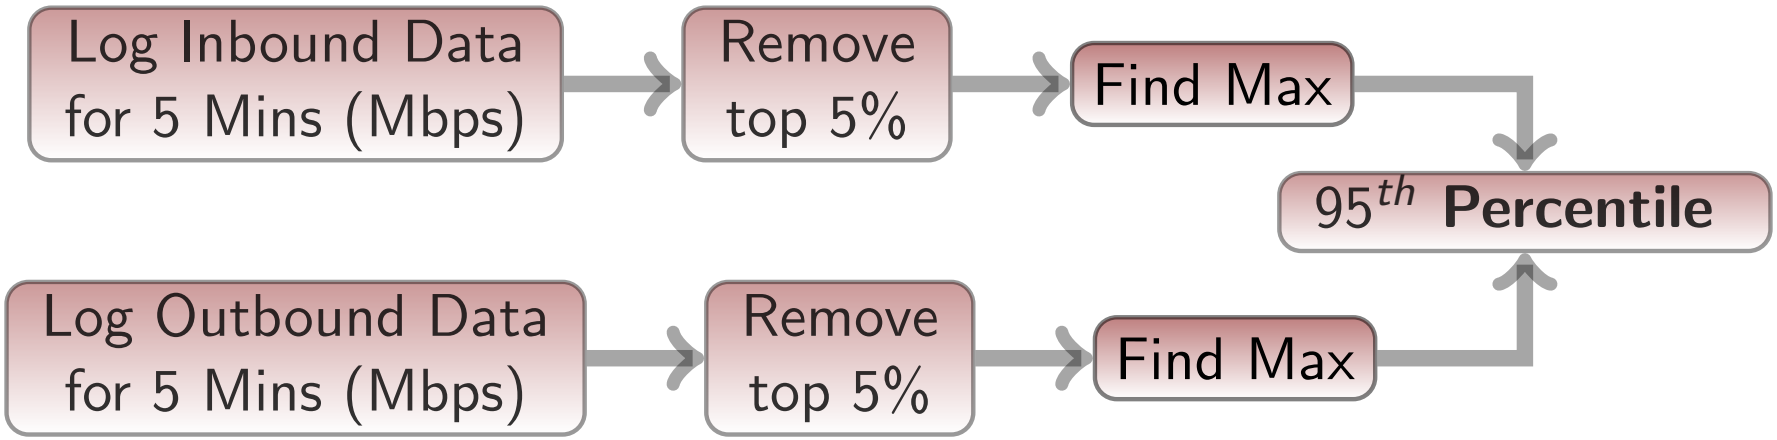

95^{th} Percentile billing

- 95^{th} Percentile billing has withstood the test of time because
  - It gives an approximation of the peak usage while ignoring some bursts
  - It is easy to implement with very small computational complexity
  - It can be calculated independently without the knowledge of other users traffic
- Although the billing mechanism has remained the same, traffic patterns have changed significantly
- For example there was a huge rise and fall of P2P traffic. More recently, there is a high surge in online video streaming traffic

Questions

- Which properties of 95^{th} Percentile are important to study and how do they change over history?
- Does 95^{th} Percentile billing volume correctly represents the amount of resources provided by transit provider?
- Is it fair to determine the billing volume for different types of customer networks using the same 95^{th} Percentile method?

Data Sets

SWITCH Data Set

- Transit services to educational institutes
- Transit data - used to calculate billing volumes of customers
- 10 years of data is analyzed from Jan 2003 - Dec 2012
- The data obtained has 5 min granularity

IXP Data Set

- Data logged from 3 Internet Exchanges in Europe SIX, BIX, ILAN
- Peering data - approximately proportional to transit data
- Data is collected for a few months from June 2013 to Aug 2013
- Data is collected in terms of mrtg graphs and then parsed

Properties

- 95^{th} Percentile to average ratio:
  - The average volume reflects the total volume of the customer network
  - 95^{th} Percentile gives an idea of peak traffic
  - If ratio is high it denotes that the customer is paying primarily for their bursts
- Skewness:
  - Skewness reveals how much a traffic distribution leans to one side of the mean
  - It gives information about the burstiness of the traffic
- We will see how these properties have evolved over time and how they are different for different classes of networks

Classifying Customer Networks

- The top 20% networks that contributed at least 75% of the total traffic are marked as Heavy Hitters and rest as Moderate Hitters

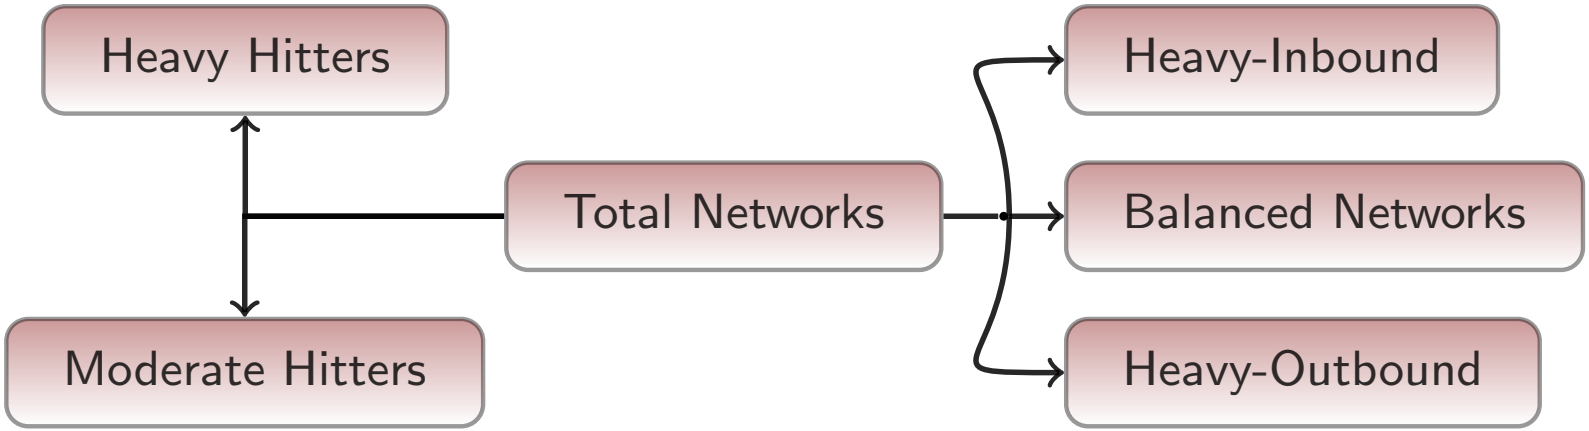

- Heavy-Inbound Networks: Inbound volume $\geq$ 2*Outbound volume
- Heavy-Outbound Networks: Outbound volume $\geq$ 2*Inbound volume
- Other networks are classified as Balanced networks

95^{th} Percentile to Average

- For the same amount of total traffic, Heavy-Inbound networks have higher 95^{th} Percentile billing volume than Heavy-Outbound networks

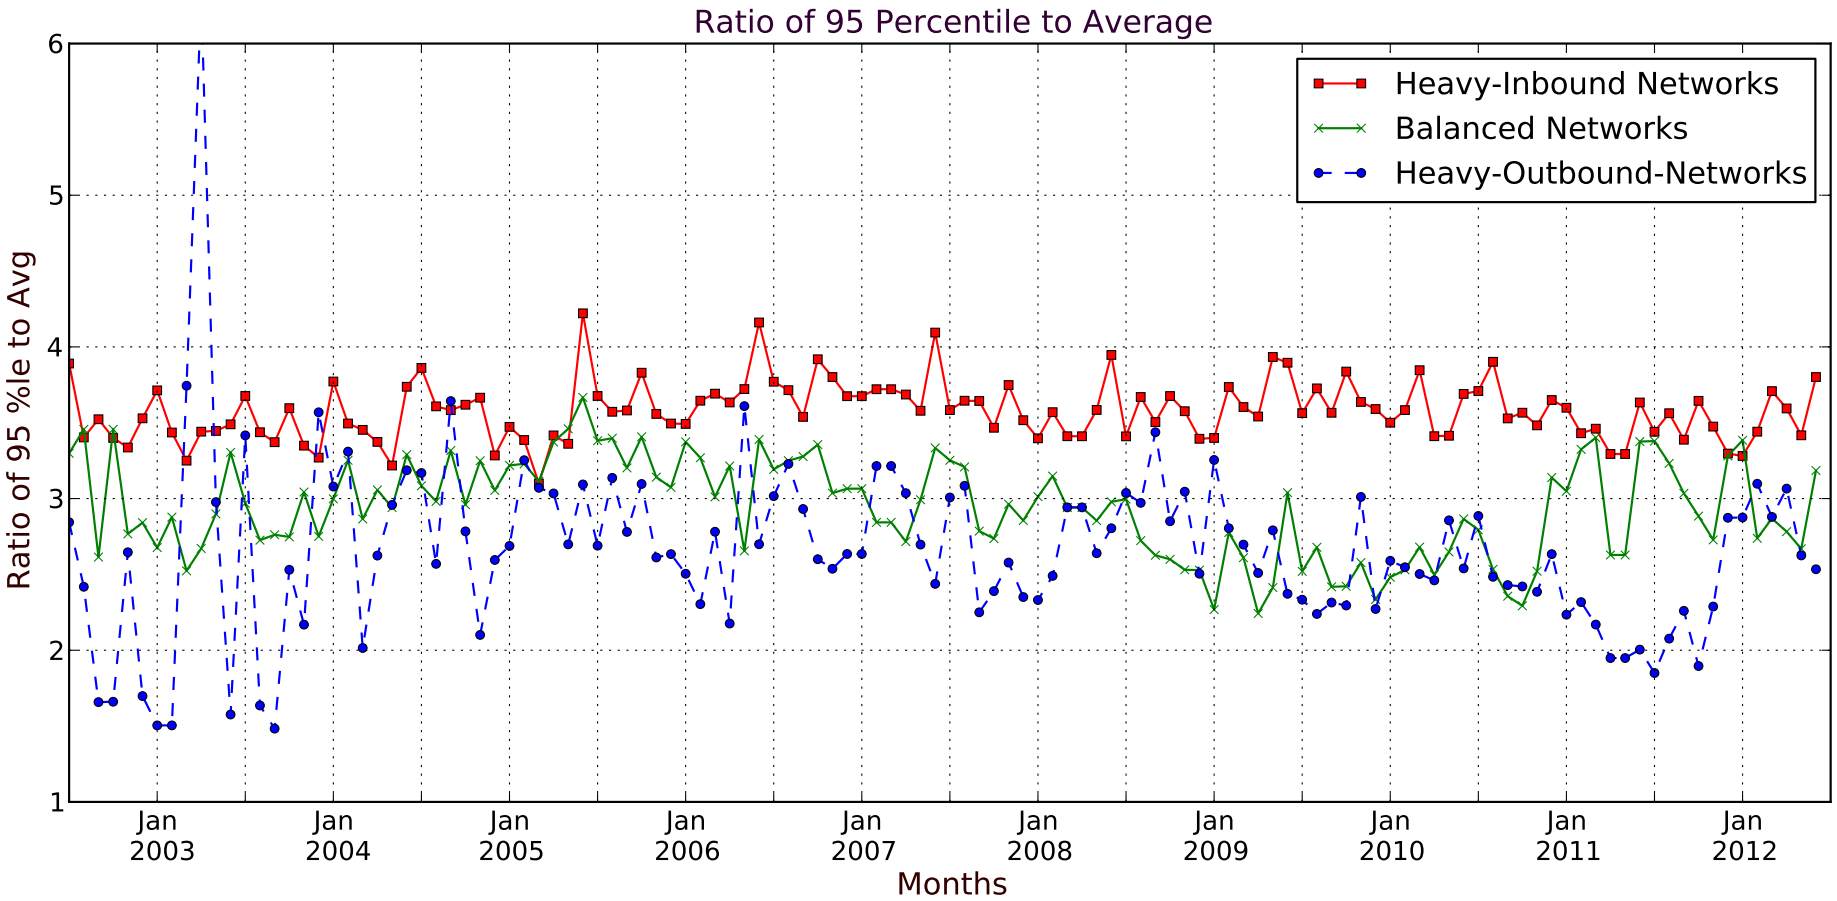

Figure: Mean 95^{th} Percentile to average ratio for different network types in the SWITCH dataset.

- Moderate hitters have higher 95^{th} Percentile to average ratio than heavy hitters

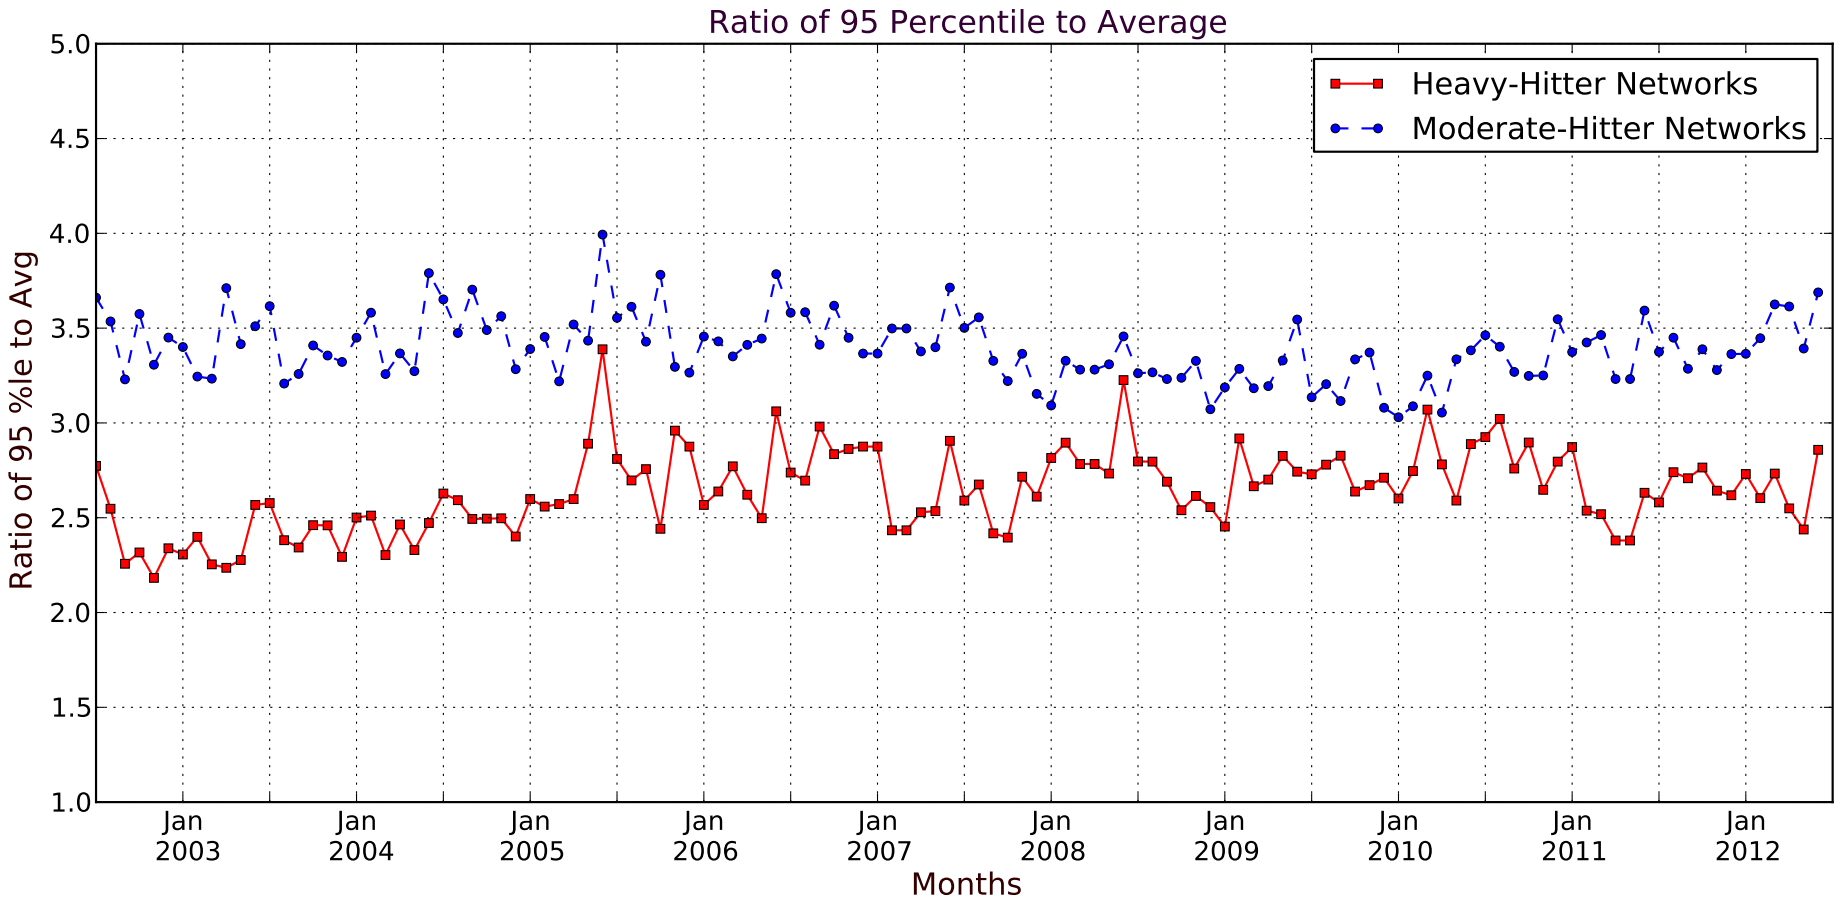

Figure: Mean 95^{th} Percentile to average ratio for different network types in the SWITCH dataset.

Skewness

- Heavy-Outbound networks have a higher skewness than Heavy Inbound networks, especially in the last 4 years

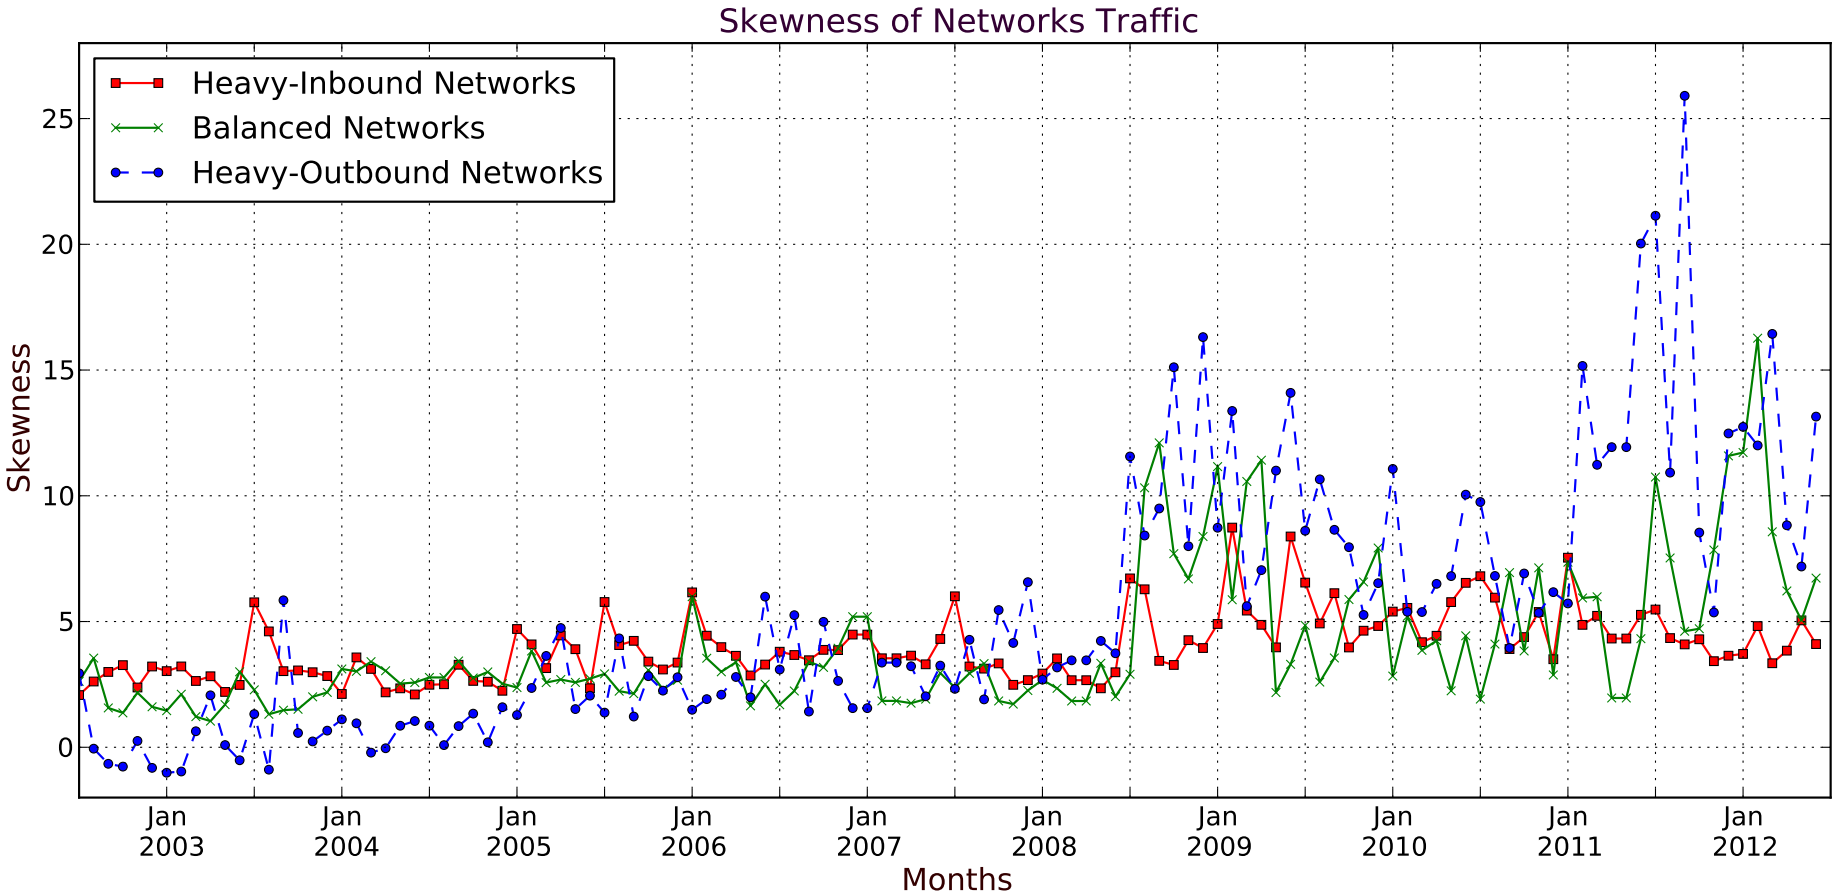

Figure: Mean skewness for different network types in the SWITCH dataset.

Skewness

- Heavy-hitter networks have larger skewness than moderate-hitter networks especially, in the last four years

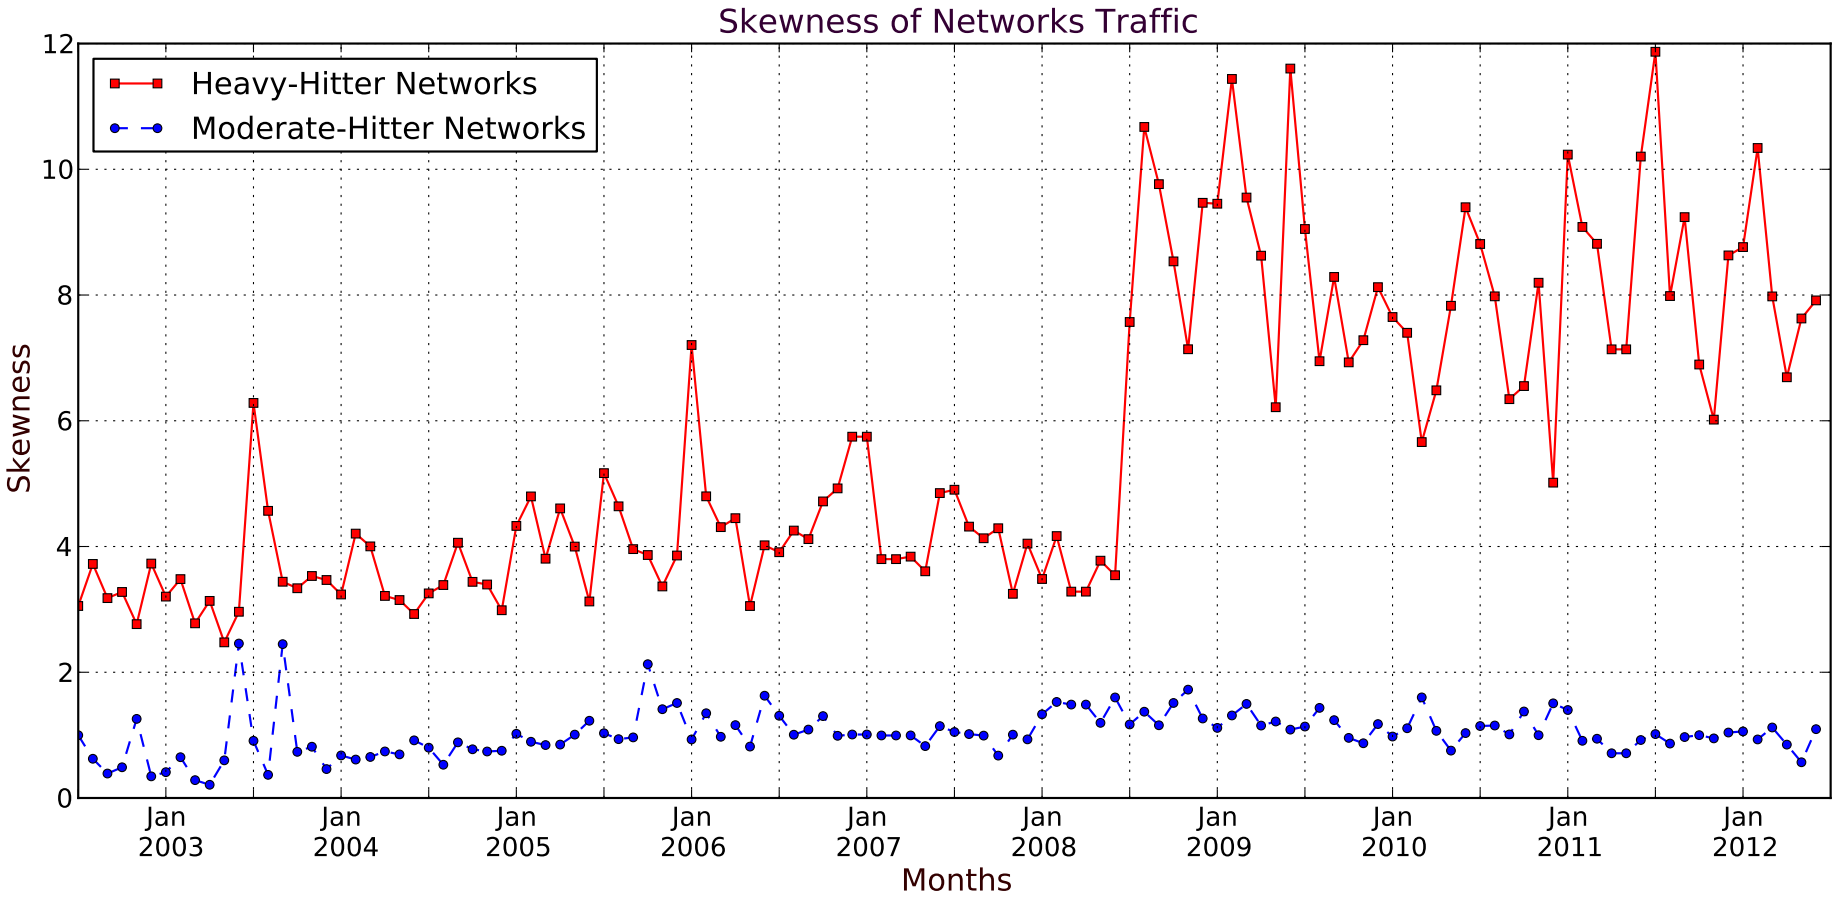

Figure: Mean skewness for Heavy and Moderate Hitter networks in the SWITCH dataset.

Observations

- The 95^{th} Percentile to average ratio more or less remained constant over the years
- Some networks like Heavy-Outbound and Heavy-Hitter have low billing volume for the same amount of traffic compared to other networks
- These networks might be employing some clever traffic shaping to reduce their bursts and in turn their transit costs

Fairness of 95^{th} Percentile

- We consider a billing mechanism fair, *if the amount of resources that a network uses is reflected in its billing volume*
- Is there a scientific way of calculating the amount of resources used? – Shapley Value
- Shapley Value gives a fair way of allocating value in a Cooperative-Game with many desired properties like efficiency, symmetry, additivity etc

Shapley Value Percentile Allocation

- We model the ISP cost allocation as a cooperative game with value function($\mathcal{V}$) of a group equal to 95^{th} Percentile of the total traffic of the group
- Then the Shapley Value of the group is uniquely defined by

$$\phi_i = \frac{1}{\mathcal{N}!} \sum_{\pi \in \Pi} \left[\mathcal{V}(S(\pi, i)) - \mathcal{V}(S(\pi, i) \backslash i) \right]$$

with i

without i

- Here, $S(\pi, i)$ is the set of all users before i in permutation π of $\mathcal{N}$ users including i
- Essentially, we calculate the difference in value of a group with and without user i. Averaging this over all possible groups gives Shapley Value of user i

Shapley Value Percentile Allocation

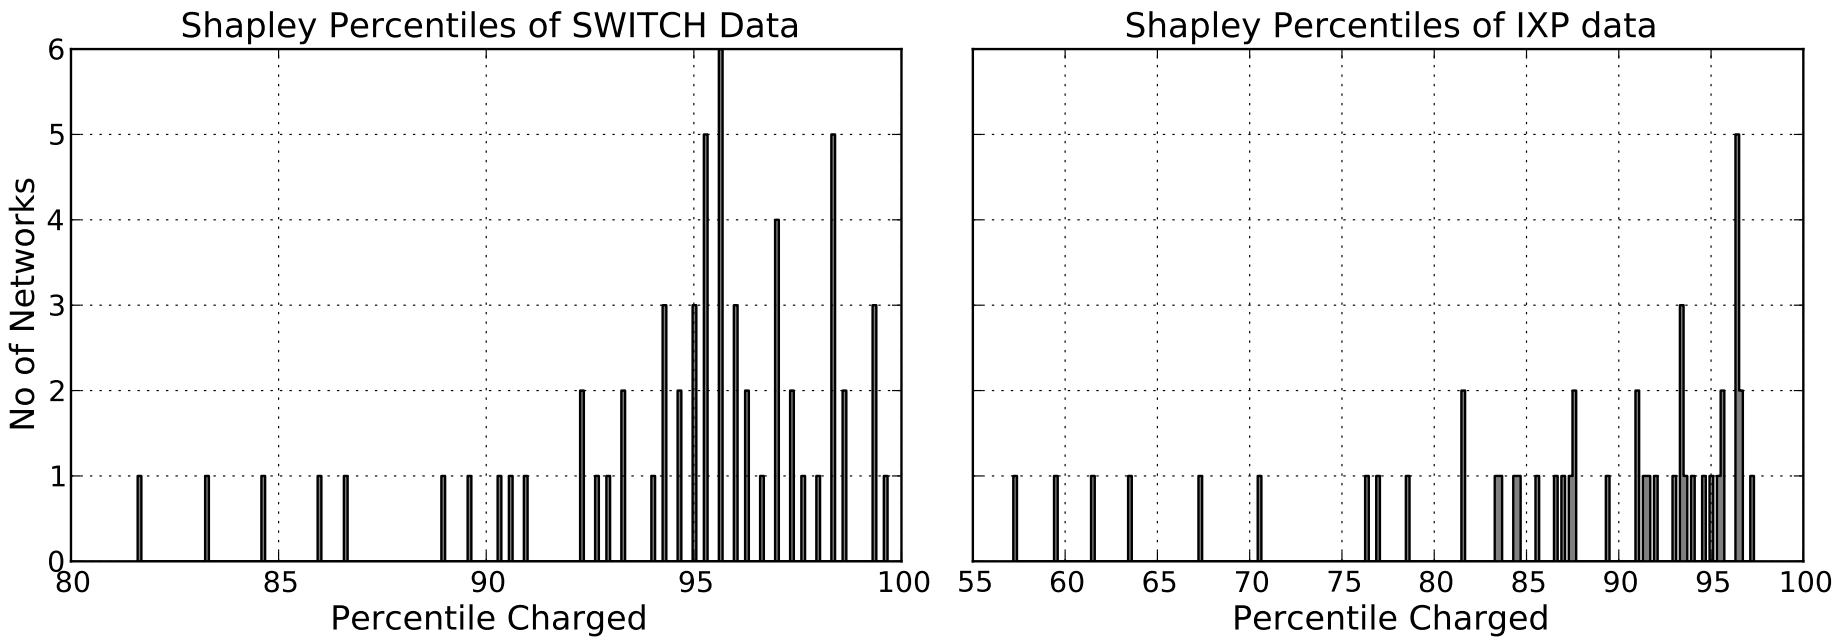

Figure: Shapley value percentiles: SWITCH dataset (Mar 2012) and IXP dataset (SIX, Aug 2013).

- Each Shapley value is converted into a percentile value to obtain the equivalent Shapley value percentile

Comparison

- Shapley value percentile bills users in a wide range from as low as 57 to as high as 98 in our examples
- 95^{th} Percentile mechanism doesn't capture this differentiation among users and charges every user the same percentile - 95^{th}
- A user that has a larger contribution during overall traffic peak has a higher Shapley value as compared to a user that has traffic during off peak periods

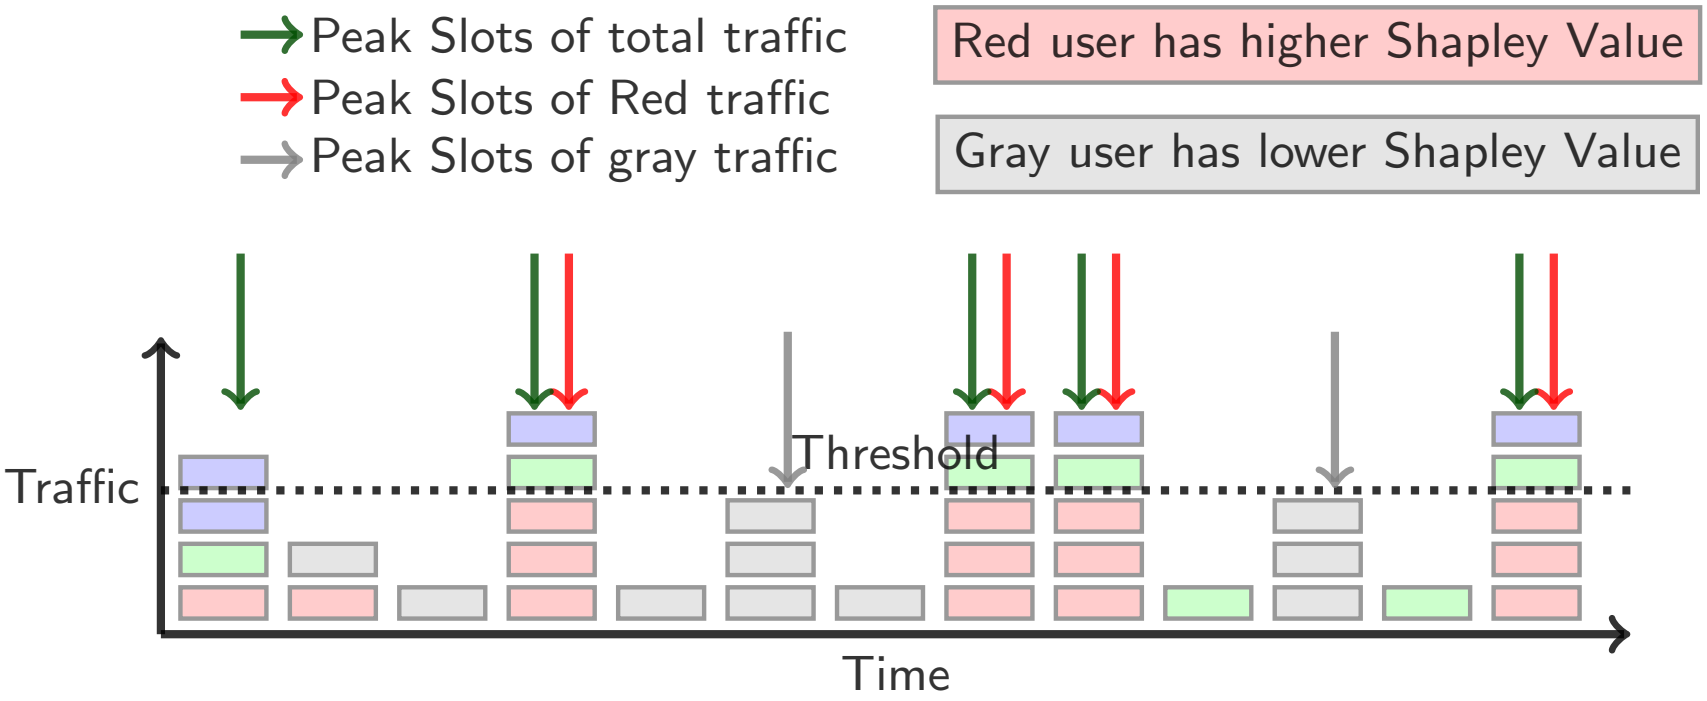

Overlap Rank vs 95^{th} Percentile rank

- **Peak Slots:** Time slots during which the total traffic exceeds a threshold
- **Overlap Rank:** Rank based on number of peak slots of the network that overlap with the peak slots of total traffic
- **Percentile Rank:** It is the rank of the network based on the 95^{th} Percentile traffic volume

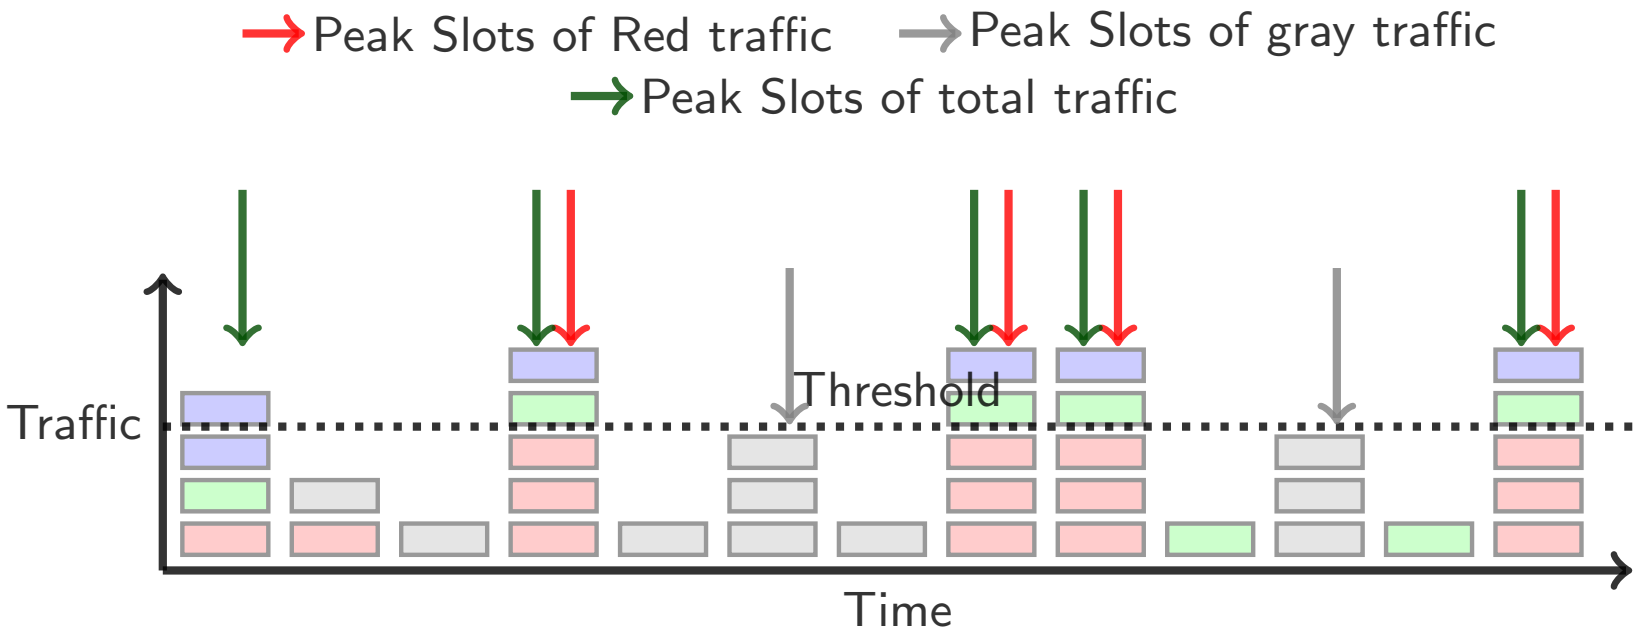

Overlap rank vs 95^{th} Percentile rank

Low 95% Rank
High Overlap rank
→ High billing volume
Small overlap with peak
→ over-charged

High 95% Rank
Low Overlap rank
→ Low billing volume
Large overlap with peak
→ under-charged

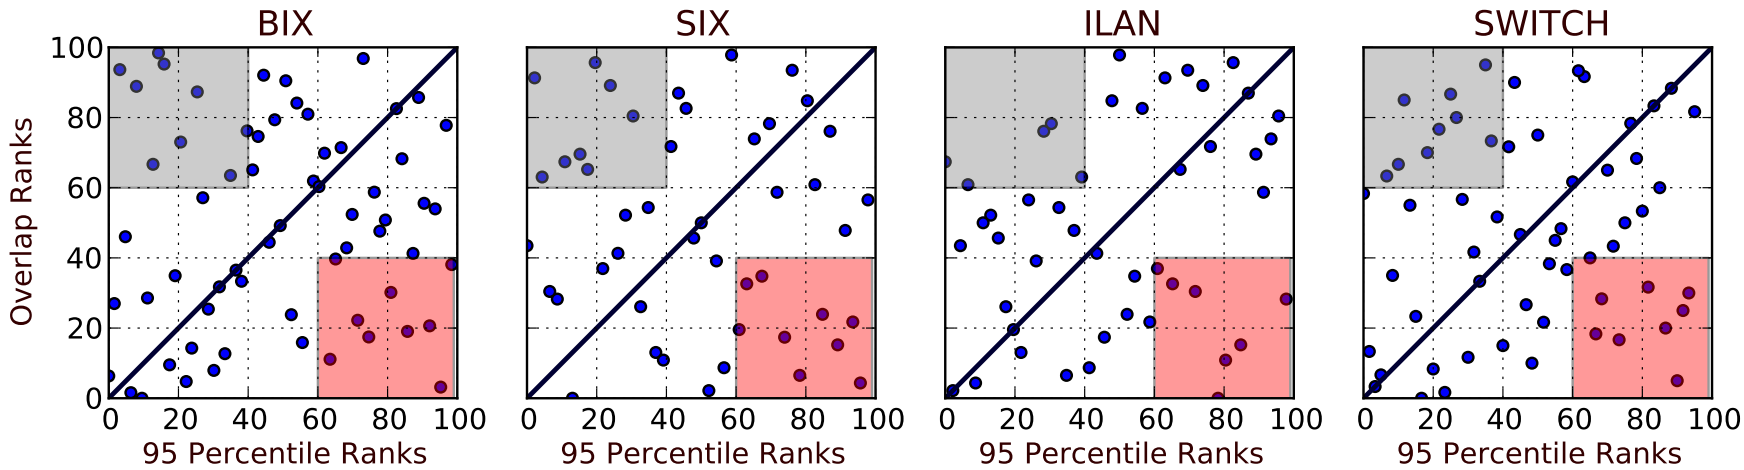

Figure: Overlap rank vs 95^{th} Percentile rank for IXP dataset (Aug 2013) and one month of SWITCH dataset (Jan 2012).

Need for a proxy

- The 95^{th} Percentile mechanism creates unfairness in charging customers without taking into account the peak traffic
- Though Shapley Value percentile is appealing it has its own shortcomings
  - Computational complexity - $\mathcal{O}(\mathcal{N}!)$. For $\mathcal{N} = 50$ complexity is 10^{64}
  - Wide range of billing percentiles anywhere between 0 to 100
  - No flexibility to provide incentives say to high volume users
- There is a need for a proxy to Shapley value which is easy to compute and also captures its essence - *Provision Ratio*

Provision Ratio

- We define Provision Ratio (PR) as

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

- It is the fraction of a network's peak traffic that occurs during total peak periods
- Though provision ratio may not be equal to the Shapley Value in an absolute sense, it has has the right characteristics in a relative sense

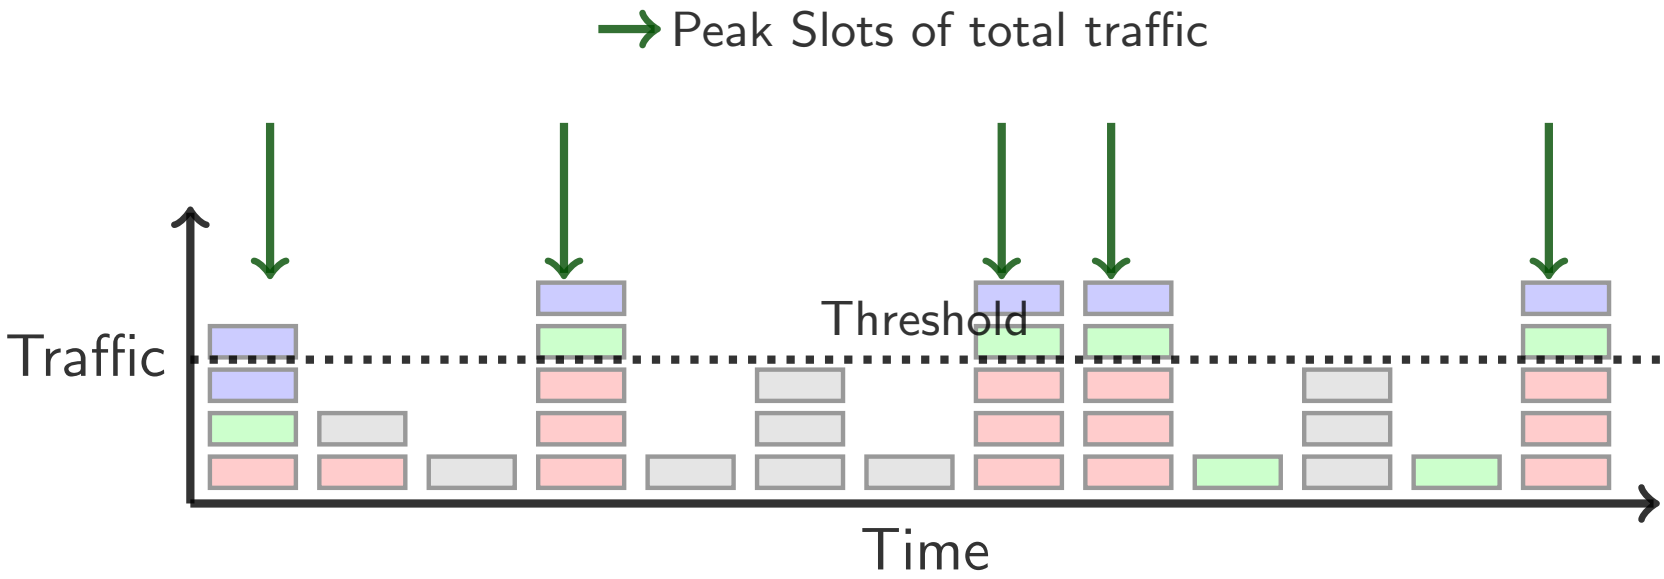

Provision Ratio

- To quantify the similarity, we consider all possible pairs of users
- For each pair we evaluate if the order is preserved
- We consider that order is preserved if

$$\text{Shapley of User 1} \geq \text{User 2} \implies \text{Provision Ratio of User 1} \geq \text{User 2}$$

- For the SWITCH and IXP data we found that the percentage of orders preserved is always above 75 %

Conclusions & Future Work

- We observed that 95^{th} Percentile to average ratio almost remained constant
- We found that charging networks based on 95^{th} Percentile billing mechanism is unfair to some users
- We defined a parameter called Provision Ratio which is easy to compute and acts as a proxy for Shapley Value
- We are developing a billing mechanism using this provision ratio which is fair and also not computationally intensive

Thank You

email id: vamseedhar.reddyvari@gmail.com
paper url: http://people.tamu.edu/~vamseedhar.reddyvaru/