CIPT: Using Tuangou to Reduce IP Transit Costs

Sergey Gorinsky

(Joint CoNEXT 2011 work with Rade Stanojevic and Ignacio Castro)

Institute IMDEA Networks, Madrid, Spain

CAIDA Workshop on Internet Economics (WIE 2011)

La Jolla, California, USA

12/2/2011

Outline

1. Internet Protocol (IP) transit costs
2. Cooperative IP Transit (CIPT)
3. Data-driven evaluation
    1. Data collection
    2. CIPT gains
4. Beyond gains sharing
5. Open problems and conclusion

Introduction: IP (Internet Protocol) transit

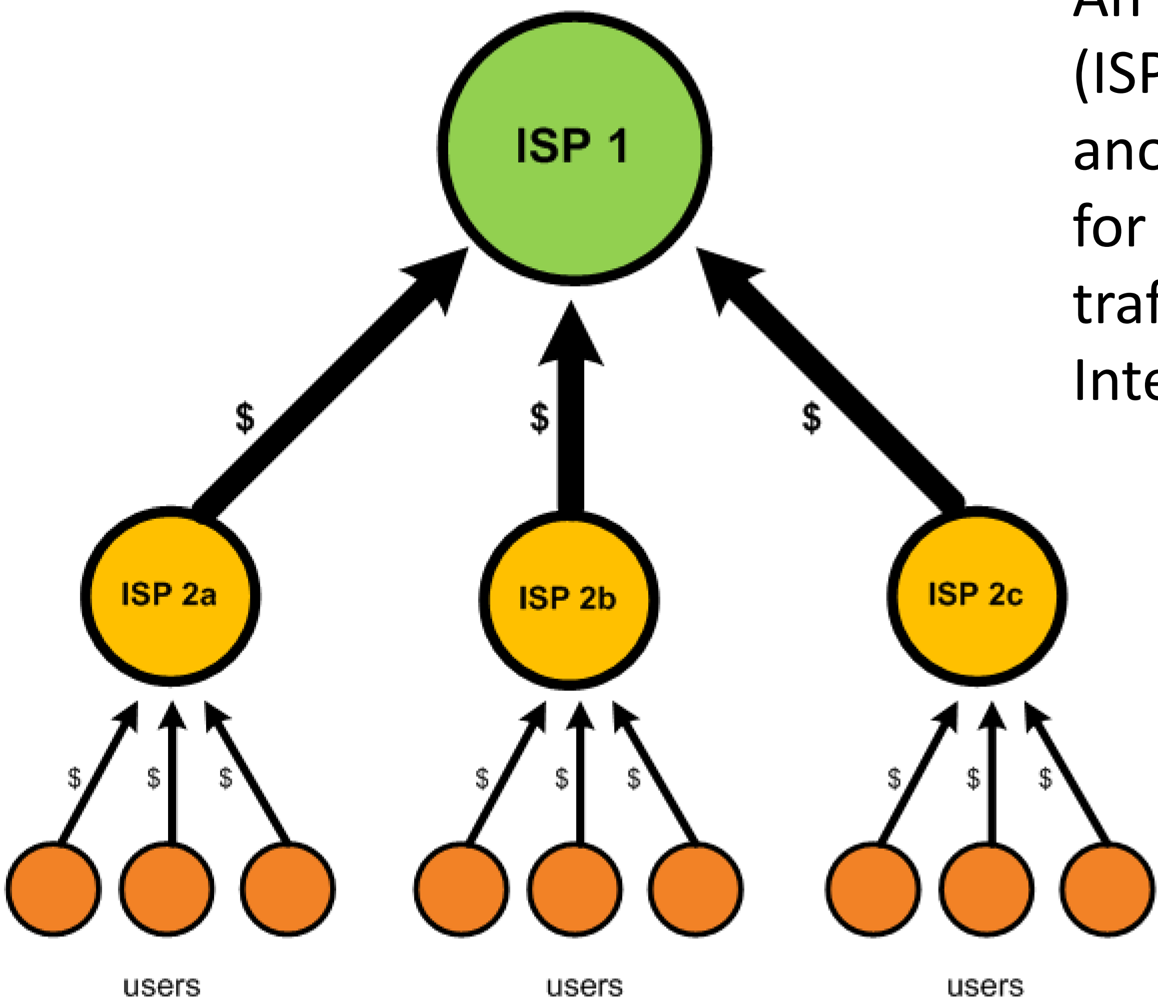

IP transit

An Internet Service Provider (ISP), the costumer, pays another ISP, the provider, for having its bidirectional traffic reaching the global Internet

IP transit billing

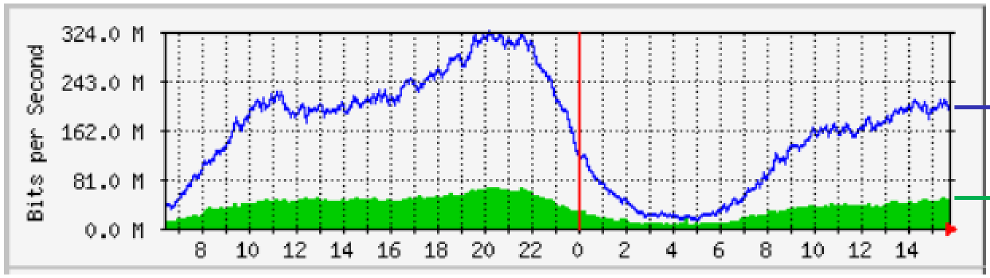

- Traffic metering
  - peak (traffic)= 95th percentile of short term traffic rates
  - SUM= peak (upstream) + peak (downstream)
  - MAX= Max [peak(upstream), peak(downstream)]

- Subadditive pricing

Committed Data Rate ($Mbps$)	Price per $Mbps$ per month
10	\$25
50	\$15
100	\$10
1000	\$5
10000	\$4

Voxel pricing
Source= https://www.voxel.net/ip-services
(as accessed on June 2011)

Motivation

Per-Mbps transit price decline vs. interdomain traffic growth

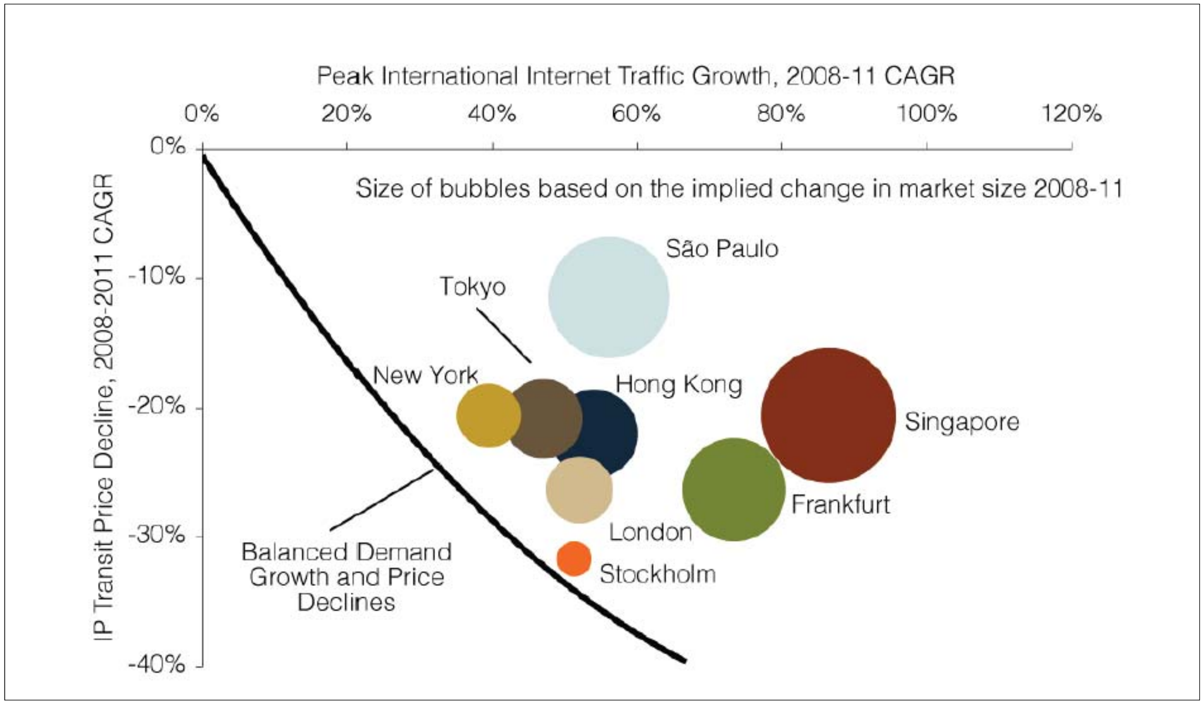

Source= https://www.telegeography

ISPs seek to reduce transit costs

Existing approaches to cost reduction

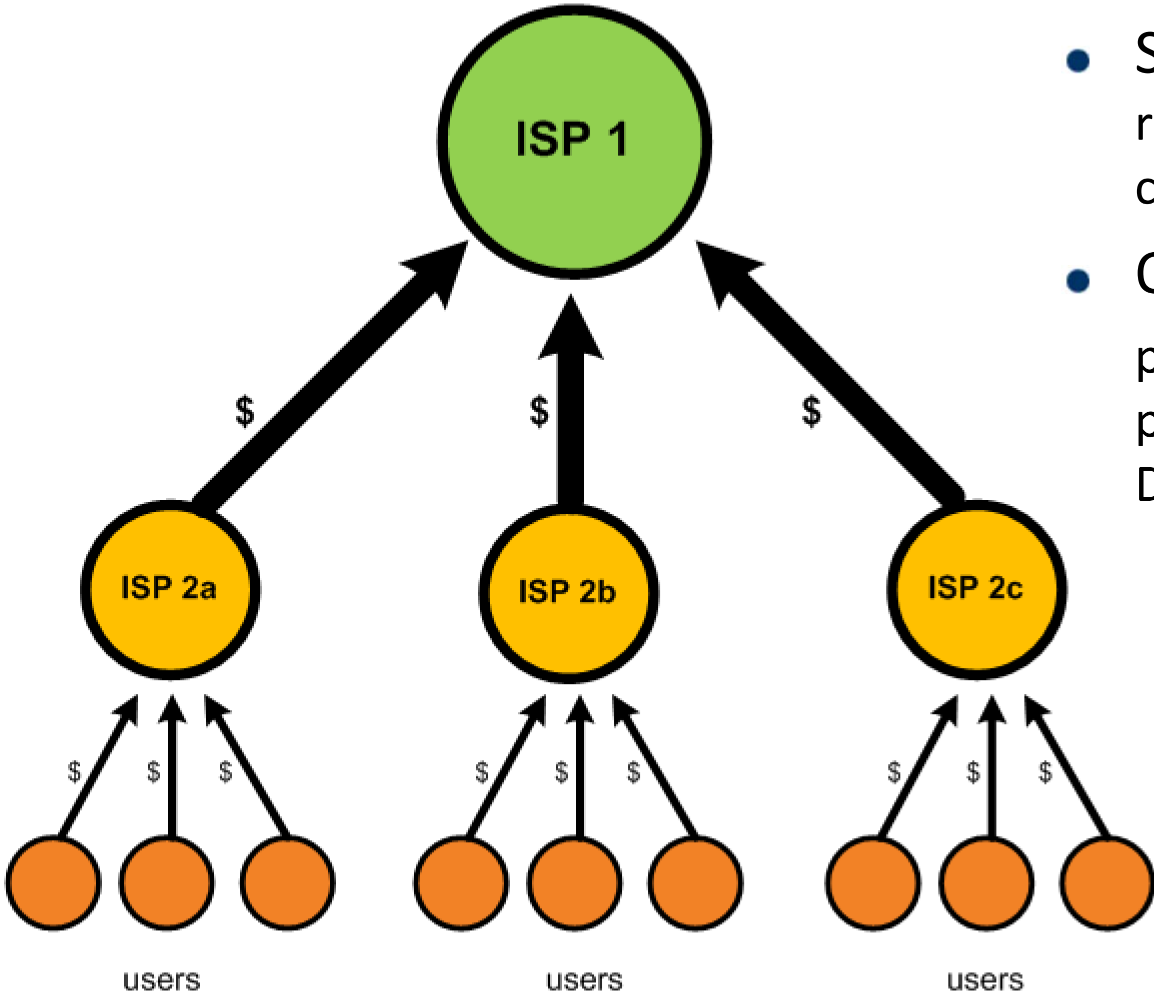

- Altering transit traffic
  - Settlement-free peering: reciprocal exchanges of own customer traffic between two ISPs
  - Other techniques: paid peering, multicast, peer-to-peer localization, Content Distribution Networks (CDNs)…

Existing approaches to cost reduction

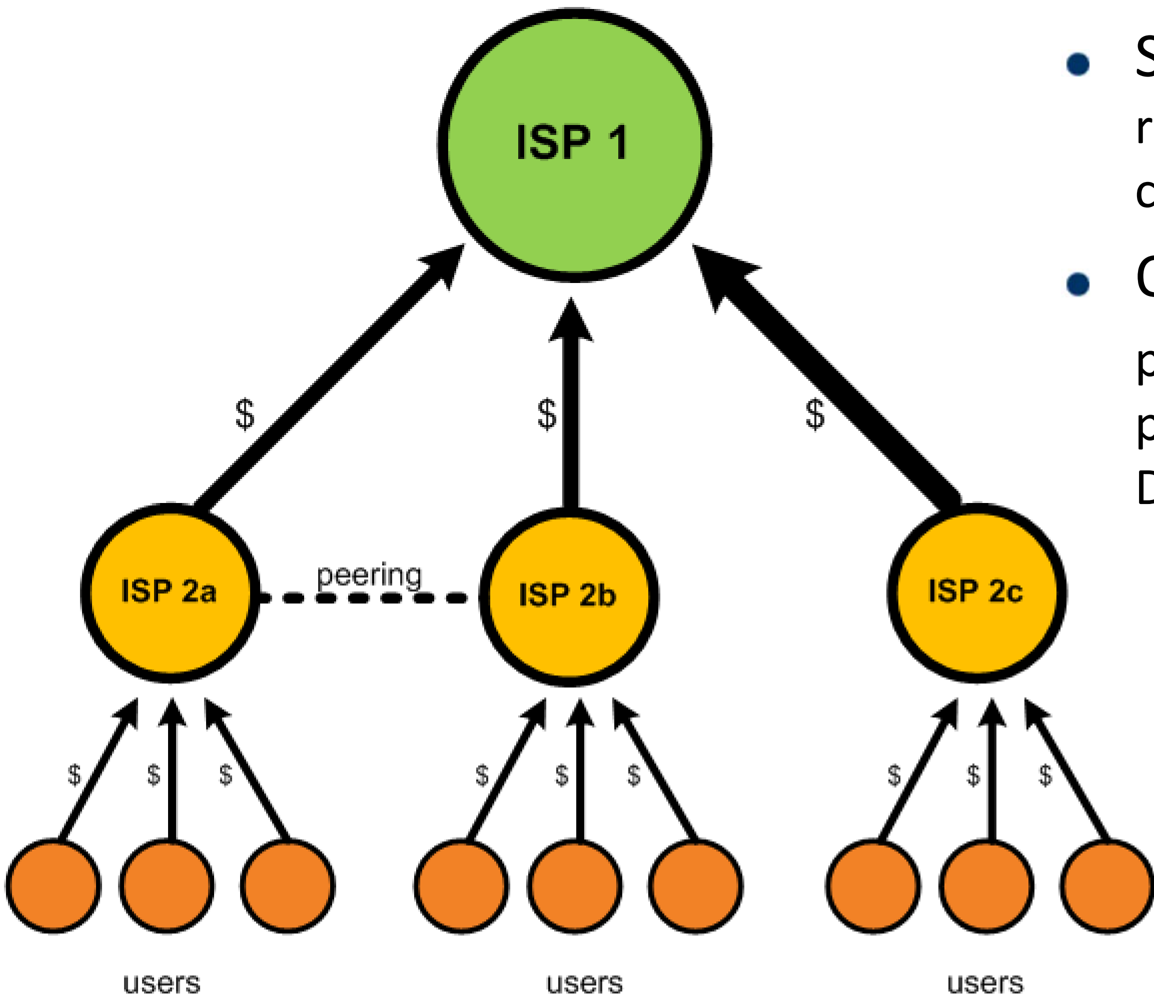

- Altering transit traffic
  - Settlement-free peering: reciprocal exchanges of own customer traffic between two ISPs
  - Other techniques: paid peering, multicast, peer-to-peer localization, Content Distribution Networks (CDNs)…

Contributions

- CIPT (Cooperative IP Transit)
  - Novel cost-reduction concept
    - Does not alter traffic
    - Reduces per-Mbps price
  - Coalitional arrangement among multiple ISPs
    - Is modeled as a cooperative game
    - Uses Shapley value to distribute gains
  - Data-driven analysis
    - Collects data from IXPs' (Internet eXchange Points) websites
    - Estimates transit traffic
    - Evaluates aggregate and individual gains

Cooperative IP Transit (CIPT) concept

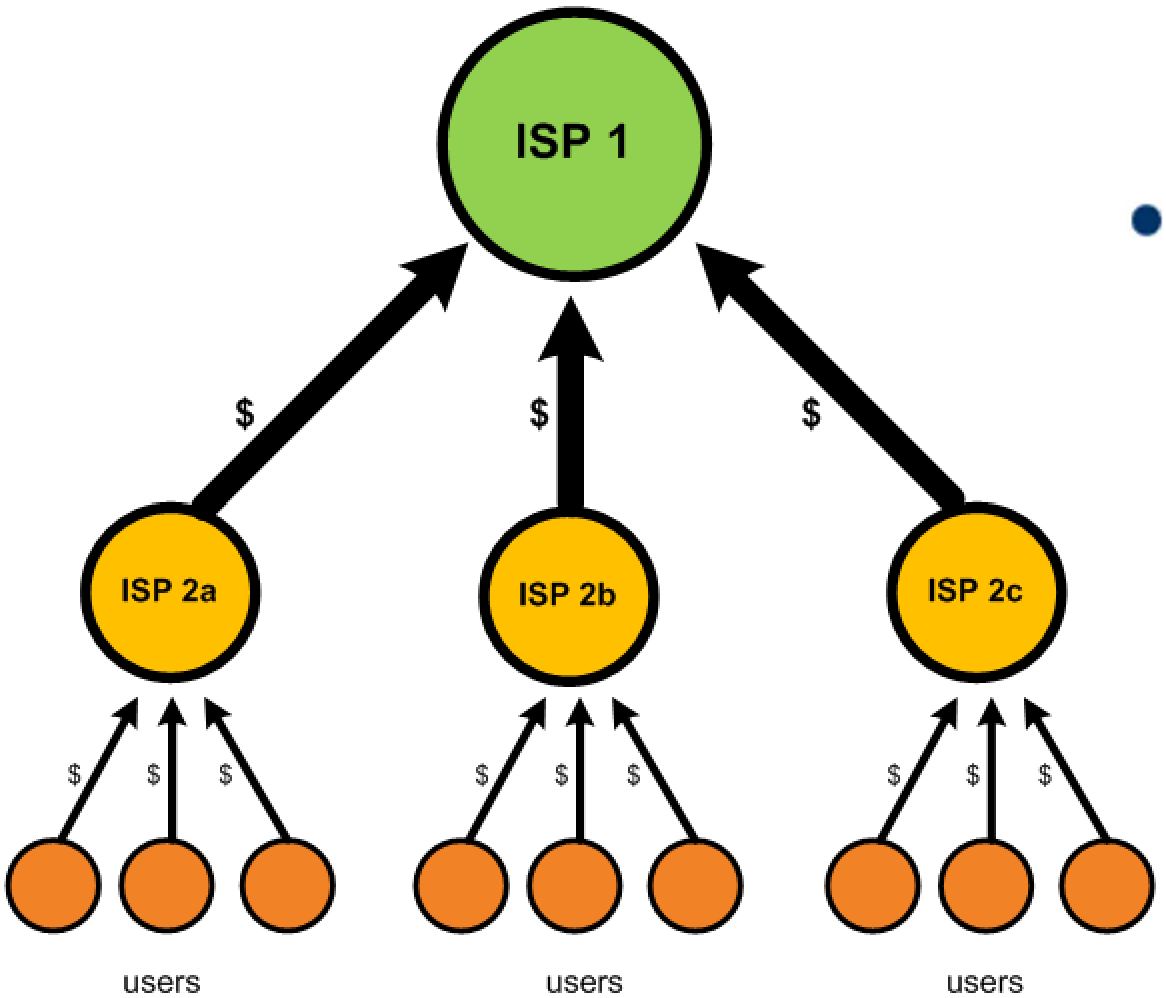

- Tuangou

  coalitional arrangement for bulk buying of IP transit

- CIPT gains

  Per-Mbps price reduction thanks to subadditive billing

Cooperative IP Transit (CIPT) concept

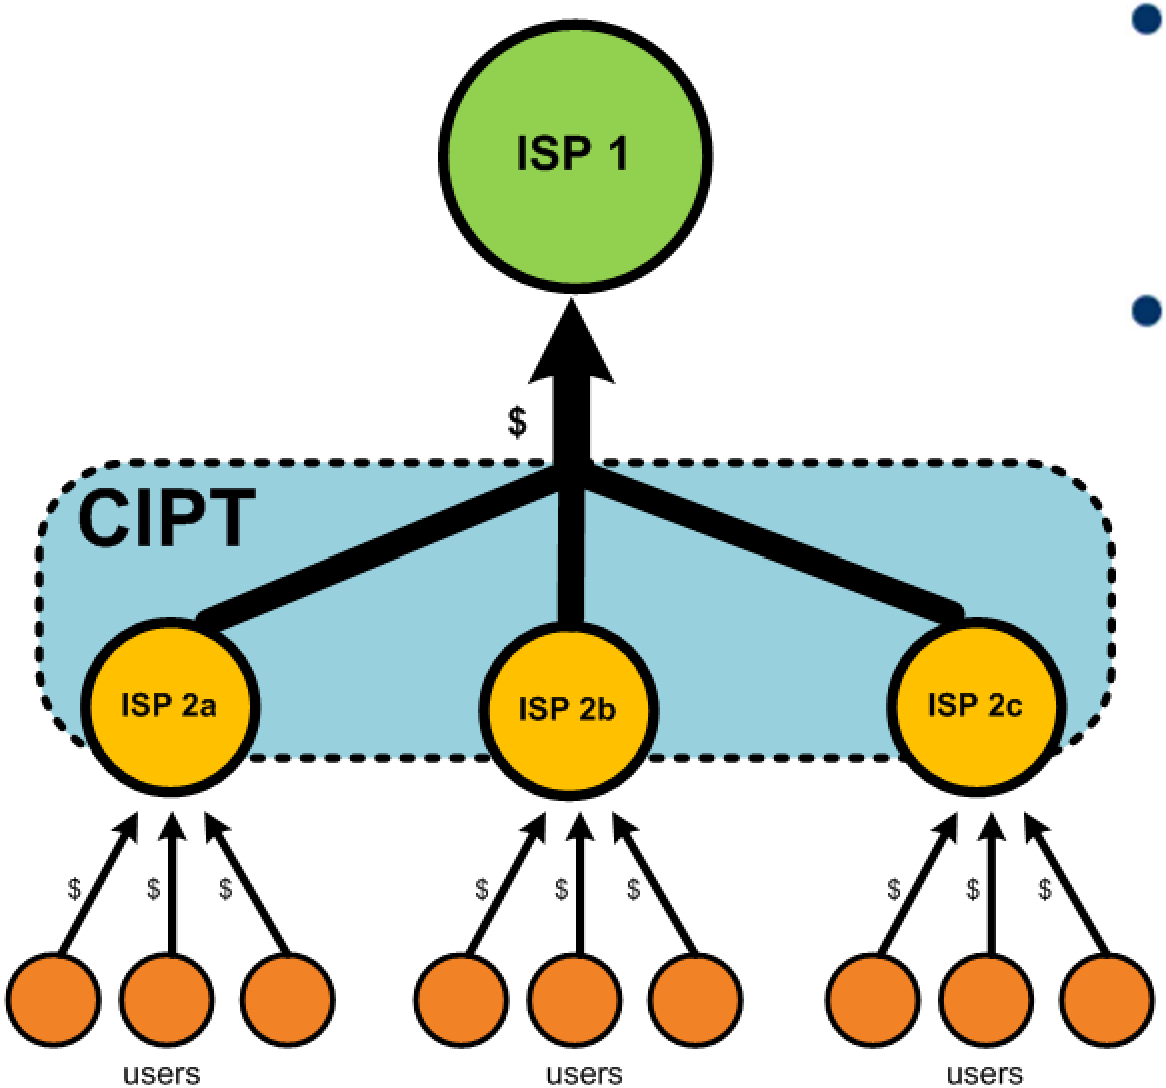

- Tuangou

  coalitional arrangement for bulk buying of IP transit

- CIPT gains

  Per-Mbps price reduction thanks to subadditive billing

Shapley value for CIPT gains' sharing

- Expected marginal contribution of a player to overall CIPT gains
- Properties
  - Existing and unique for any cooperative game
  - Fair, efficient, symmetric, additive and null-player
  - Individually rational
- Calculation
  - Hard to calculate exactly
  - Estimated accurately by our Monte-Carlo method

Shapley value definition

- Shapley value(i)

ISP *i*'s expected marginal contribution if the players join the coalition one at a time, in a uniformly random order

$$\phi_i(c) = \frac{1}{N!} \sum_{\pi \in S_N} \underbrace{\left(c(S(\pi, i)) - c(S(\pi, i) \setminus i)\right)}_{i\text{'s marginal contribution}}$$

N = number of players
c(S)= cost of coalition *S*
S(π, i) = set of players arrived in the system not later than *i*
π = permutations of the set of players *N*

Shapley value estimation

- Monte Carlo method*
  - We estimate the Shapley value as the average cost contribution over set π_k of K randomly sampled arrival orders.

$$\hat{\phi}_i(c) = \frac{1}{K} \sum_{\pi \in \Pi_K} \left(c(S(\pi, i)) - c(S(\pi, i) \setminus i) \right)$$

- Estimation accuracy
  - K is the knob controlling the accuracy
  - We use K = 1000 to keep the error under 1%

(*) D. Liben-Nowell, A. Sharp, T. Wexler, K. Woods, "Computing Shapley Value in Cooperative Supermodular Games", Preprint, 2010.
R. Stanojevic, N. Laoutaris, P. Rodriguez, "On Economic Heavy Hitters: Shapley Value Analysis of the 95th-Percentile Pricing", Proc. of ACM IMC 2010.

Data-driven evaluation

1. Crawling the Internet to **collect traffic images** from IXP's websites

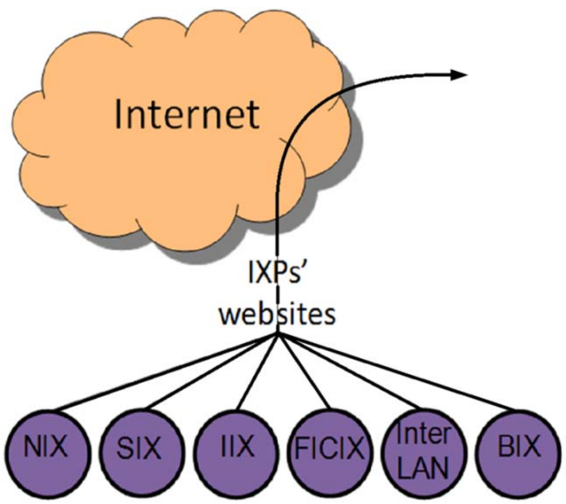

IXP	acronim	# of members	peak (*Gbps*)	average (*Gbps*)
Neutral IX (Prague)	NIX	54	116	76
Slovak IX	SIX	52	42	23
Israeli IX	IIX	17	2.1	1.38
Finnish IX	FICIX	25	32	19
InterLAN (Bucharest)	InterLAN	63	22	11
Budapest IX	BIX	53	152	92

Data-driven evaluation

1. Crawling the Internet to **collect traffic images** from IXP's websites

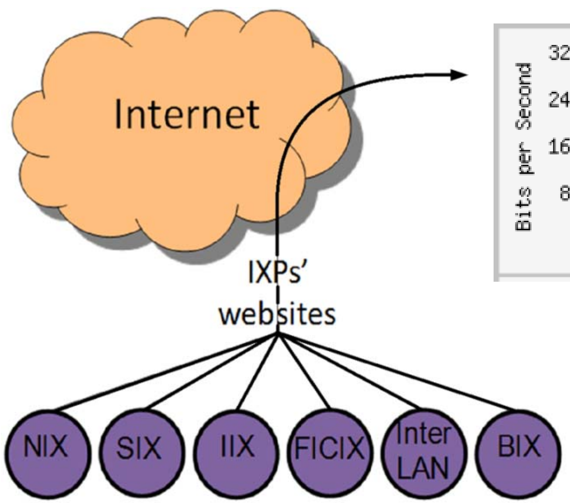

Collection of mrtg images

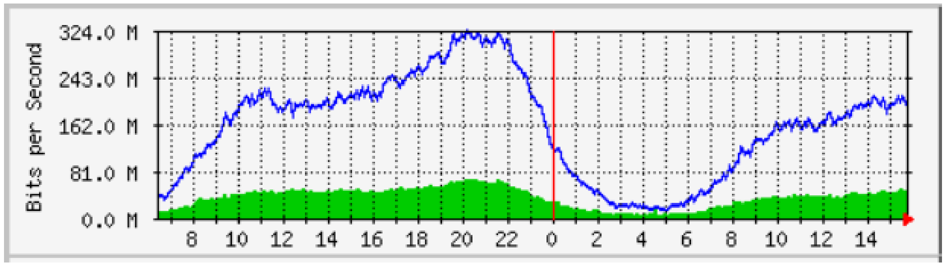

IXP	acronim	# of members	peak ($Gbps$)	average ($Gbps$)
Neutral IX (Prague)	NIX	54	116	76
Slovak IX	SIX	52	42	23
Israeli IX	IIX	17	2.1	1.38
Finnish IX	FICIX	25	32	19
InterLAN (Bucharest)	InterLAN	63	22	11
Budapest IX	BIX	53	152	92

Data-driven evaluation

1. Crawling the Internet to **collect traffic images** from IXP's websites

Collection of mrtg images

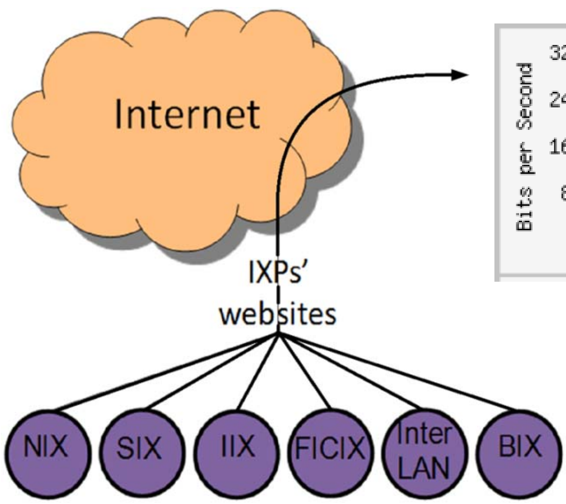

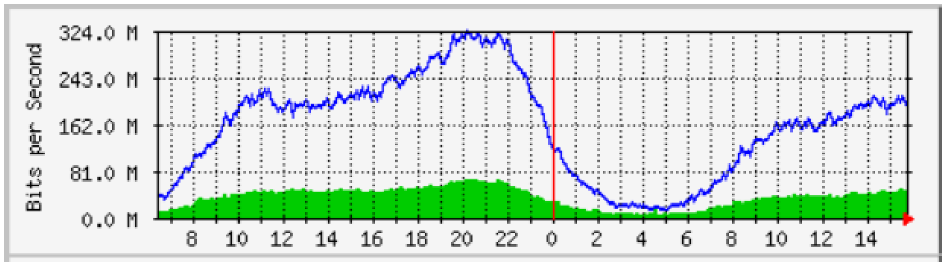

Optical Character Recognition (OCR)

IXP	acronim	# of members	peak (*Gbps*)	average (*Gbps*)
Neutral IX (Prague)	NIX	54	116	76
Slovak IX	SIX	52	42	23
Israeli IX	IIX	17	2.1	1.38
Finnish IX	FICIX	25	32	19
InterLAN (Bucharest)	InterLAN	63	22	11
Budapest IX	BIX	53	152	92

2. Transform images into **numeric data on peering traffic**

From peering to transit traffic

- Transit traffic is rarely available
- Our hypothesis is that transit traffic and peering traffic are similar
- We validate the similarity with public data from two ISPs (HEATNET and SANET)

Similarity quantification

Cosine-similarity ≈ $sim(X, Y) = \frac{\sum_{i=1}^{T} X_i Y_i}{\sqrt{\sum_{i=1}^{T} X_i^2} \sqrt{\sum_{i=1}^{T} Y_i^2}}.$

- If sim(X; Y) = 1, then X = α·Y
- Otherwise sim(X; Y) < 1

ISP	$sim(T_{up}, P_{up})$	$sim(T_{down}, P_{down})$
HEANET	0.988	0.965
SANET	0.996	0.991

Peering-transit traffic similarity

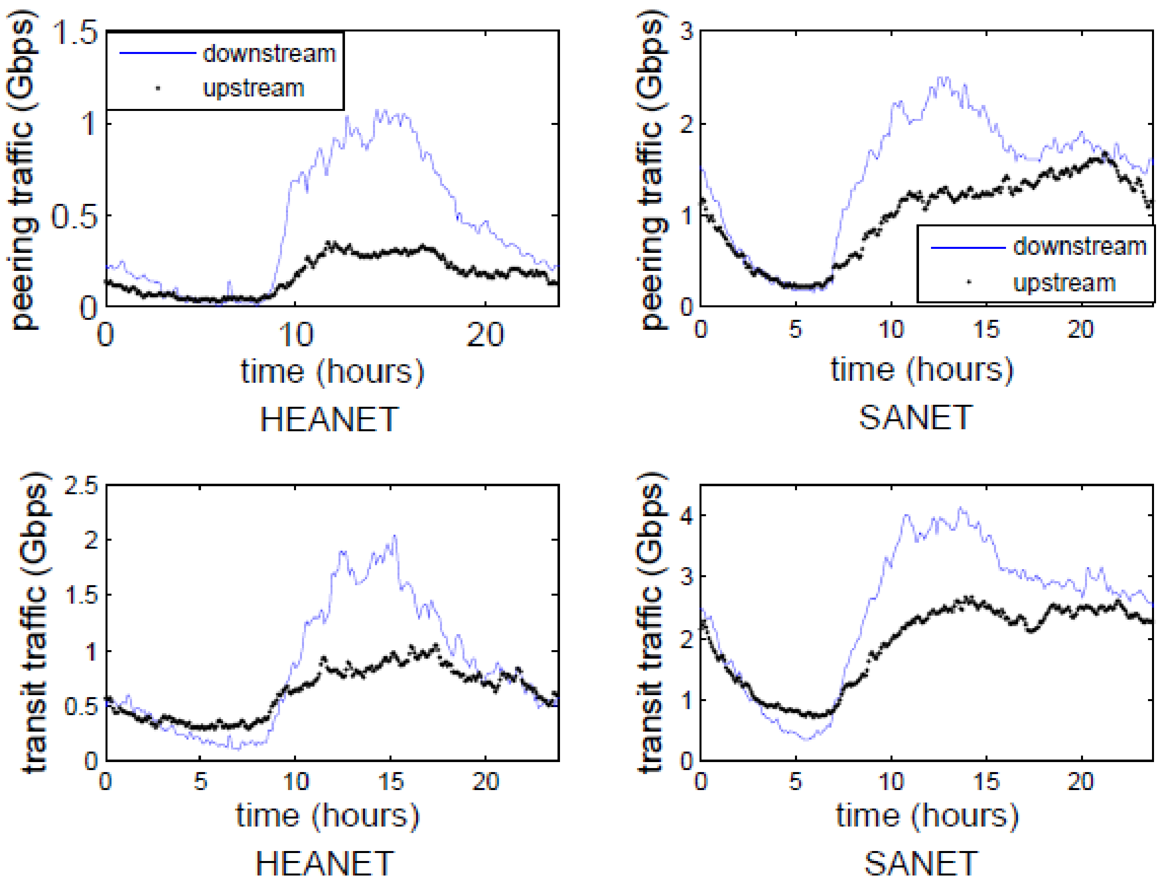

- Peering and transit follow very similar patterns
- $\alpha = 1.5$
- We scale peering traffic by α within the range [0.5;4]

Aggregate CIPT gains

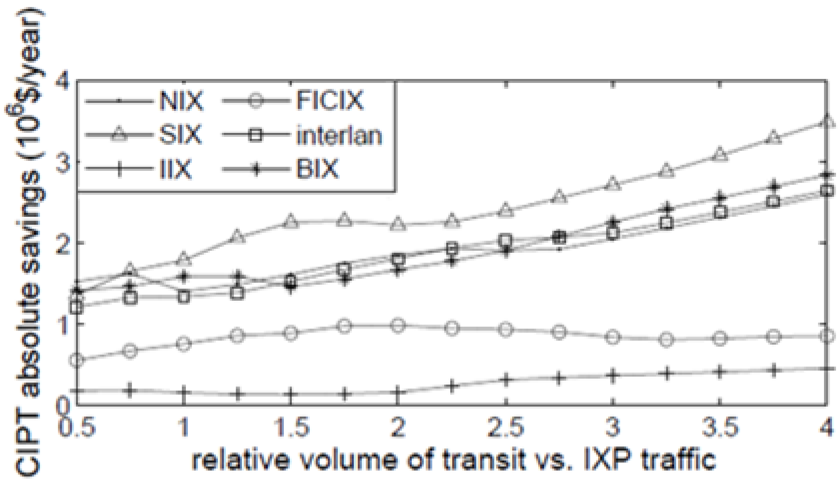

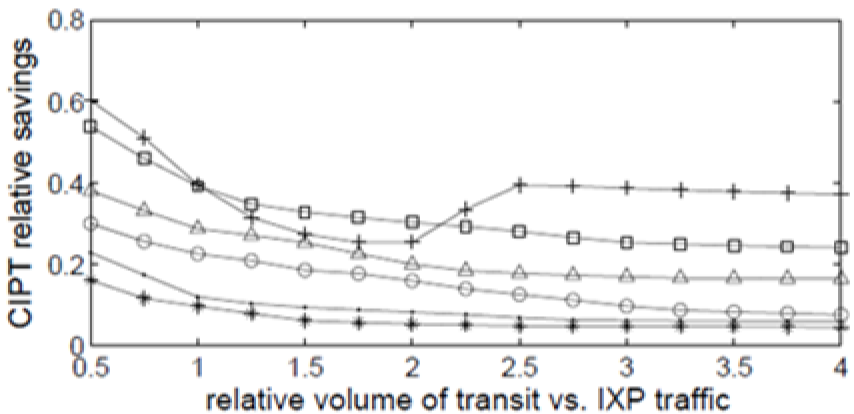

- Absolute aggregate CIPT gains grow with IXP size (in terms of billed traffic)
- Relative aggregate CIPT gains decrease with IXP size

Aggregate CIPT gains

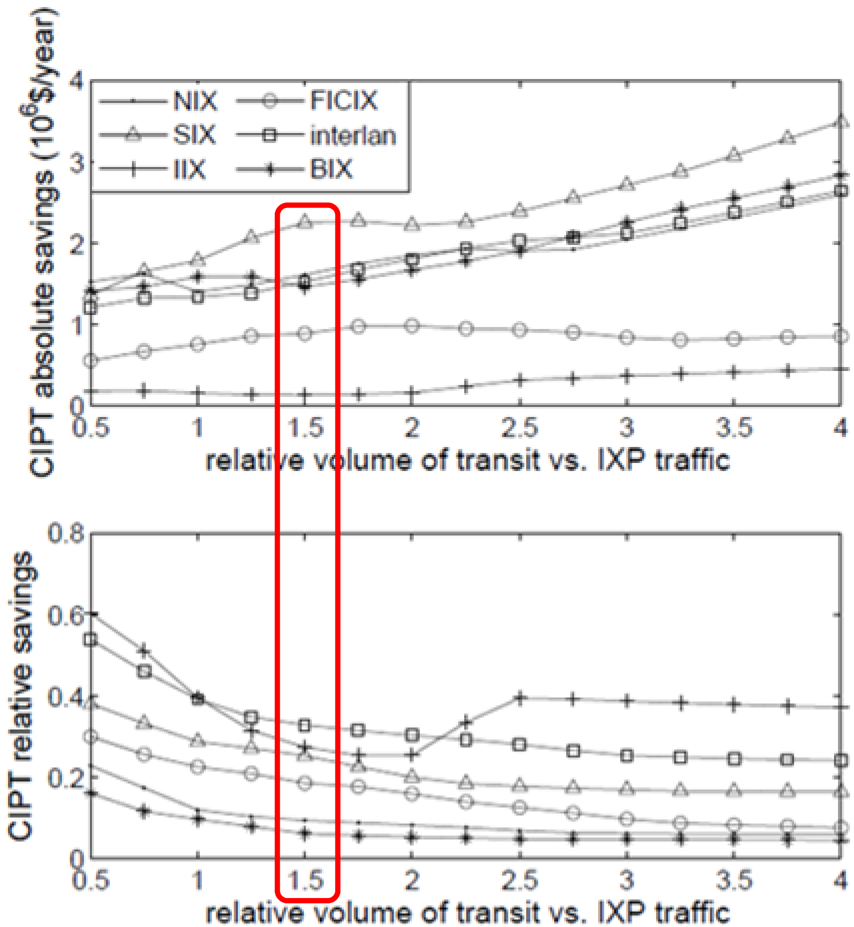

- Absolute aggregate CIPT gains grow with IXP size (in terms of billed traffic)
- Relative aggregate CIPT gains decrease with IXP size

Per-partner CIPT gains

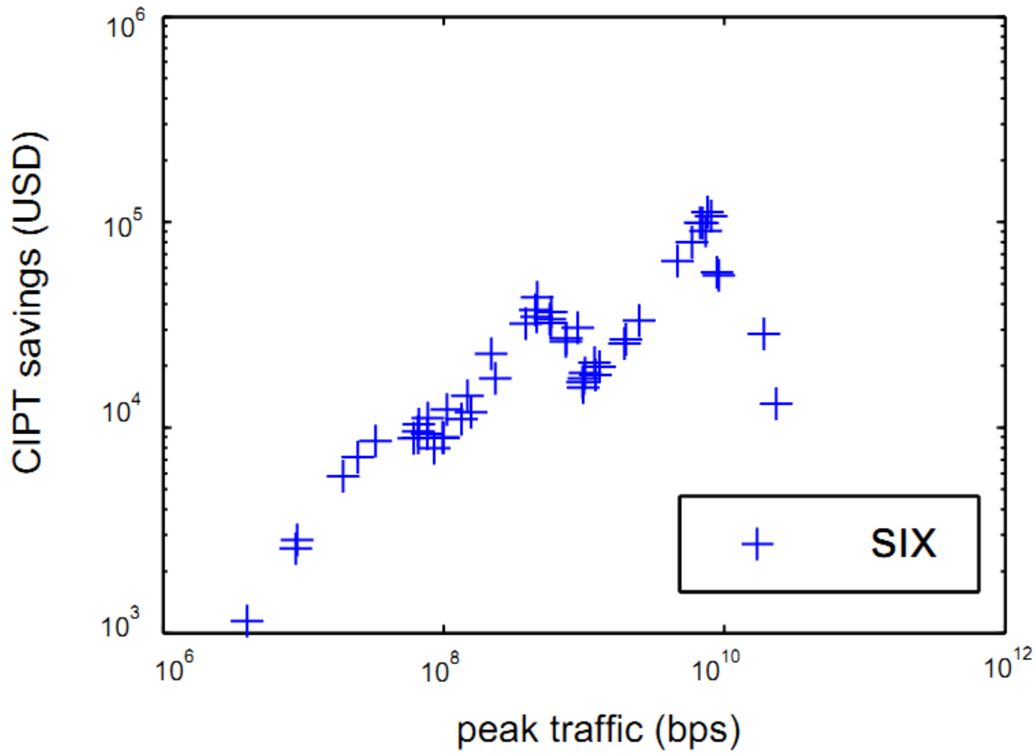

- Absolute individual CIPT gains grow with ISP size
- Relative individual CIPT gains decrease with ISP size

There are large gains for all CIPT members

CIPT gains and coalition size

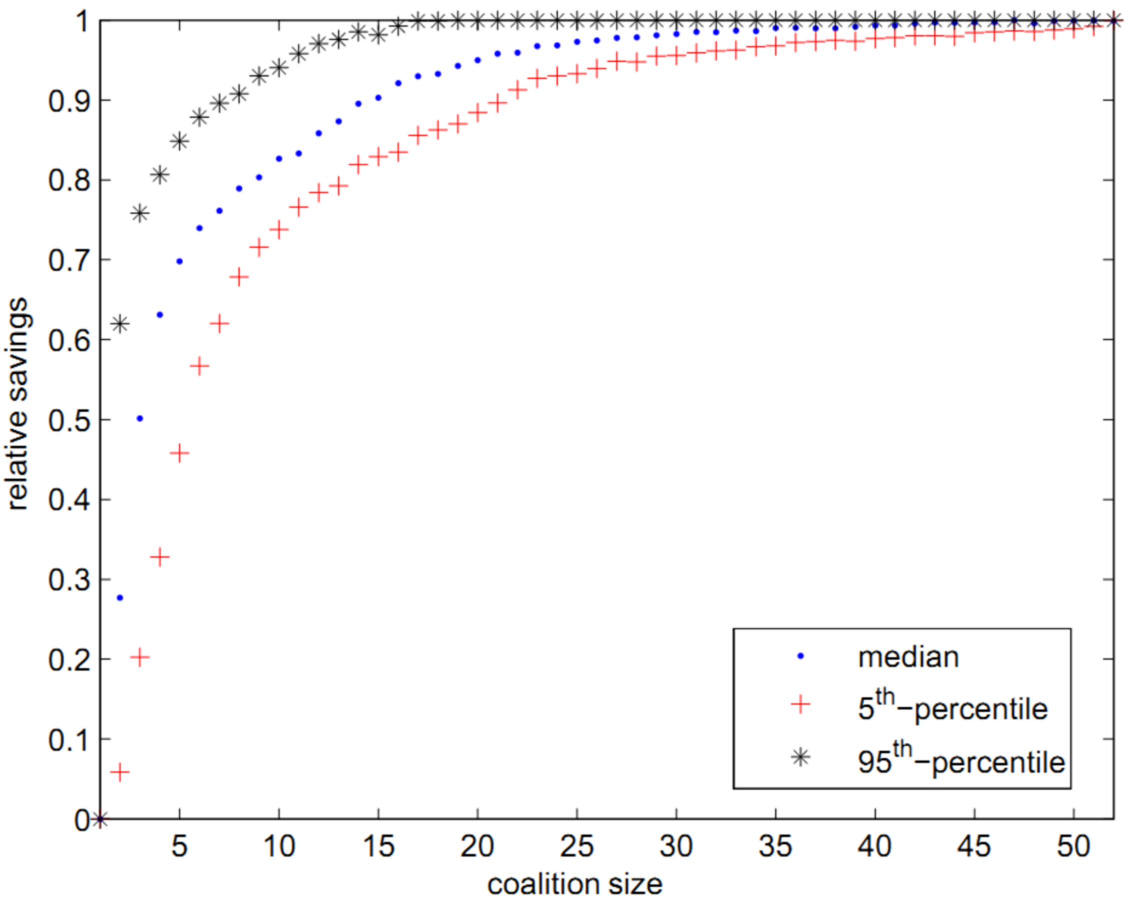

- Fraction of CIPT gains in SIX (52 ISPs)

CIPT gains and coalition size

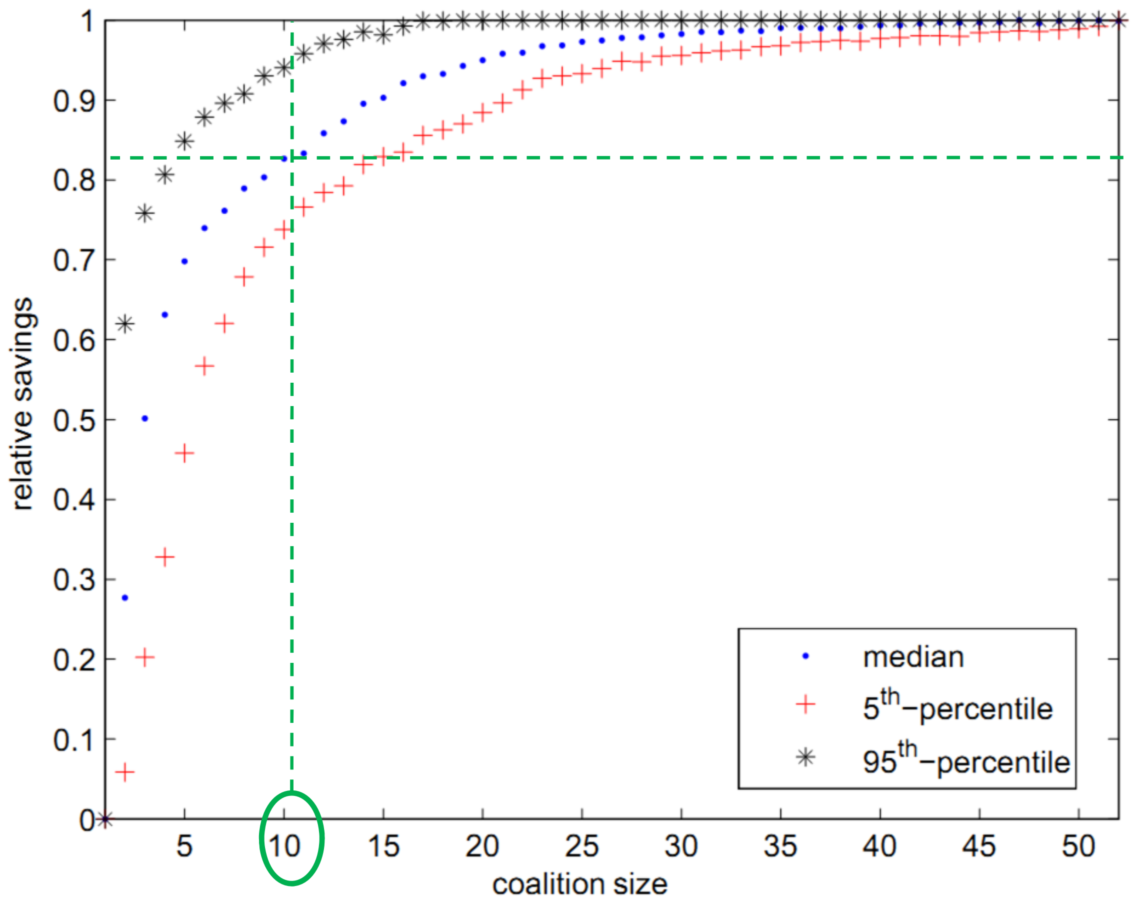

- Fraction of CIPT gains in SIX (52 ISPs)

Small coalitions provide most of total attainable gains

Beyond gains sharing

- Organizational embodiment
- Physical infrastructure
- Inter-domain routing
- Performance
- Traffic confidentiality
- Transit providers and strategic issues

Strategic issues

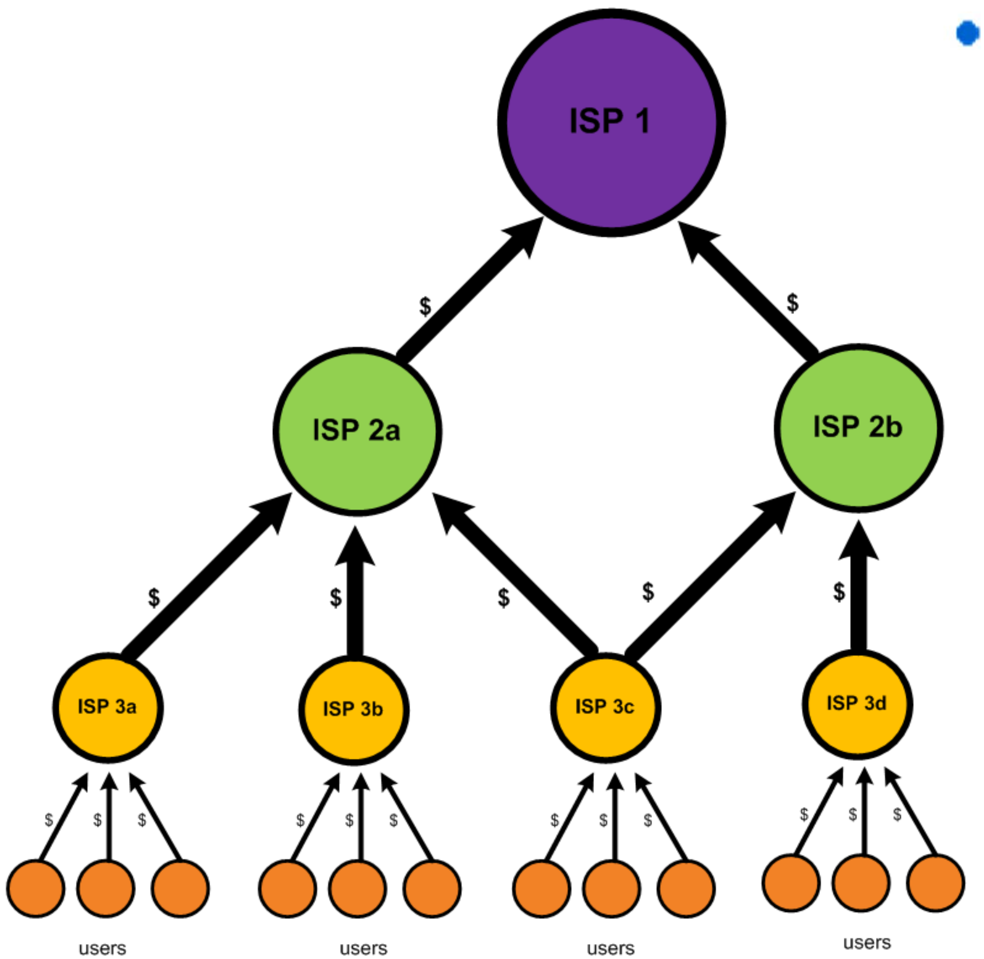

- Costs saved by CIPT coalitions are not necessarily the revenues lost by a transit provider
  - CIPT as a new customer

Strategic issues

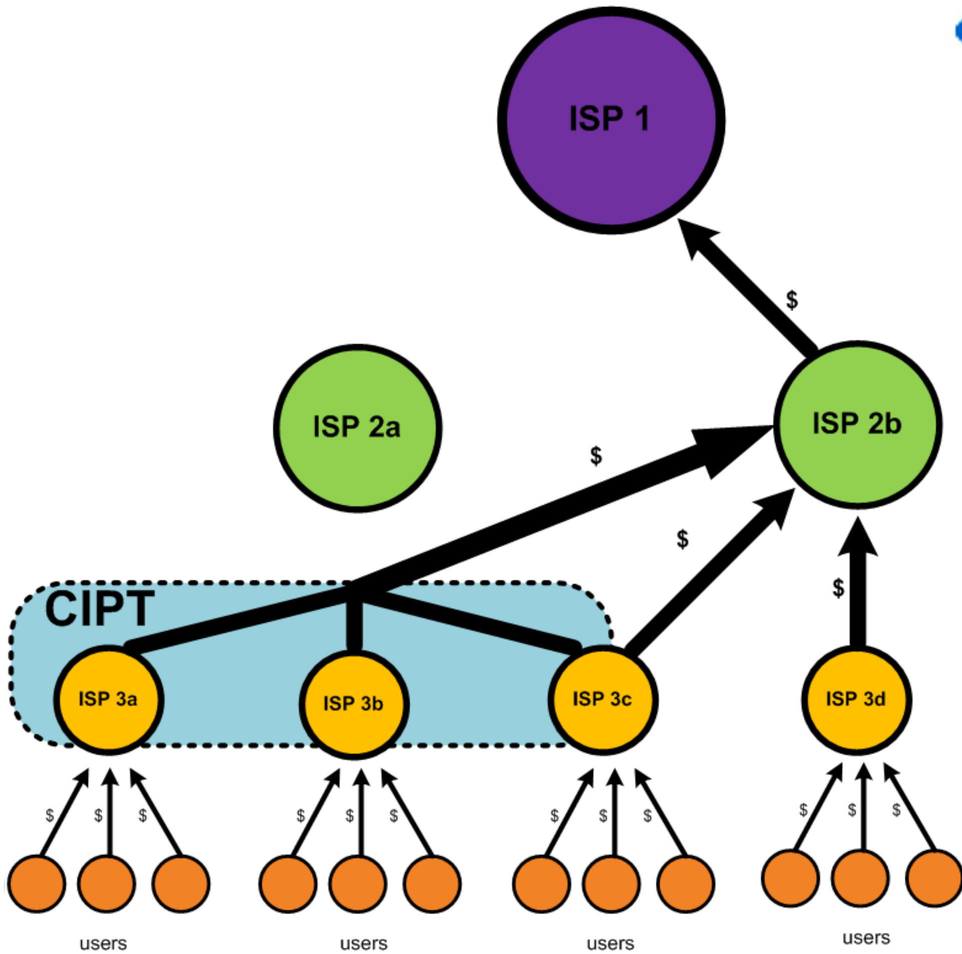

- Costs saved by CIPT coalitions are not necessarily the revenues lost by a transit provider
  - CIPT as a new customer
  - bypass the middle-man

Strategic issues

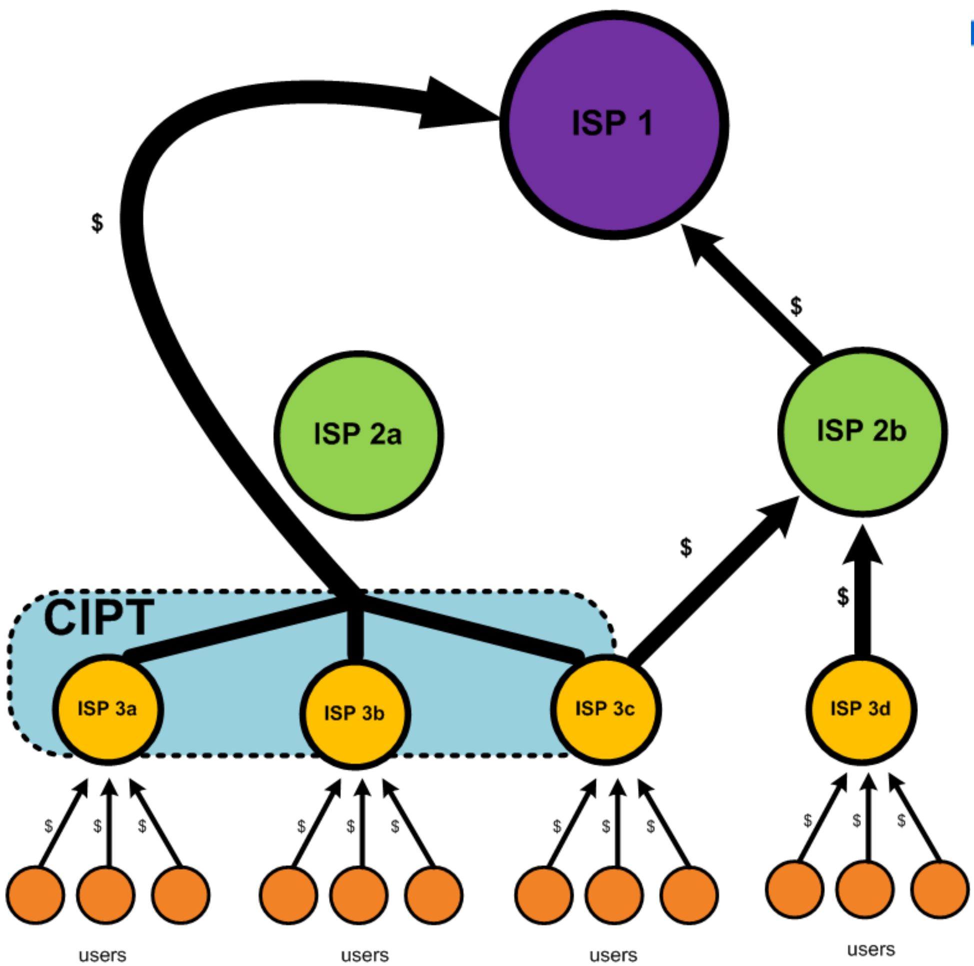

- Costs saved by CIPT coalitions are not necessarily the revenues lost by a transit provider
  - CIPT as a new customer
  - bypass the middle-man

Open problems

- **#1:** How do changes in CIPT affect its dynamic?
- **#2:** Can we quantify the factors that influence the CIPT coalition formation process?
- **#3:** Can we derive more suitable metrics that would approximate the Shapley value closely, while being explicit and simple to calculate?
- **#4:** What would be the effect of CIPT on the Internet AS-level topology?

Conclusions

- We propose a novel mechanism for IP transit cost reduction: Cooperative IP Transit (CIPT)
- CIPT reduces costs significantly through bulk buying of IP transit
- We model CIPT as a cooperative game and use Shapley value as a mechanism for cost sharing
- The evaluation of CIPT with real data shows
  - Significant aggregate and individual gains
  - Large gains even with small coalitions