technical report

TR-IMDEA-Networks-2011-4

CIPT: Using Tuangou to Reduce IP Transit Costs

Rade Stanojevic

Ignacio Castro

Sergey Gorinsky

October 2011

CIPT: Using Tuangou to Reduce IP Transit Costs*

Rade Stanojevic
Institute IMDEA Networks, Madrid

Ignacio Castro
Institute IMDEA Networks, Madrid

Sergey Gorinsky
Institute IMDEA Networks, Madrid

ABSTRACT

A majority of ISPs (Internet Service Providers) support connectivity to the entire Internet by transiting their traffic via other providers. Although the transit prices per $Mbps$ decline steadily, the overall transit costs of these ISPs remain high or even increase, due to the traffic growth. The discontent of the ISPs with the high transit costs has yielded notable innovations such as peering, content distribution networks, multicast, and peer-to-peer localization. While the above solutions tackle the problem by reducing the transit traffic, this paper explores a novel approach that reduces the transit costs without altering the traffic. In the proposed CIPT (Cooperative IP Transit), multiple ISPs cooperate to jointly purchase IP (Internet Protocol) transit in bulk. The aggregate transit costs decrease due to the economies-of-scale effect of typical subadditive pricing as well as burstable billing: not all ISPs transit their peak traffic during the same period. To distribute the aggregate savings among the CIPT partners, we propose Shapley-value sharing of the CIPT transit costs. Using public data about IP traffic and transit prices, we quantitatively evaluate CIPT and show that significant savings can be achieved, both in relative and absolute terms. We also discuss the organizational embodiment, relationship with transit providers, traffic confidentiality, and other aspects of CIPT.

General Terms

economics

*Tuangou (pronounced "twangoo"), a term originating in China, loosely translates as team buying or group buying, http://en.wikipedia.org/wiki/Tuangou.

ACM CoNEXT 2011, December 6–9 2011, Tokyo, Japan.

Keywords

network economics, IP transit, burstable billing, group-buying, cost sharing.

1. INTRODUCTION

The Internet ecosystem involves thousands of ISPs (Internet Service Providers) linked in a more or less hierarchical manner to support universal connectivity of Internet users. Only a handful of huge ISPs can access the entire Internet without paying anyone for the reachability. For the vast majority of the other ISPs, the universal connectivity comes at the price of IP (Internet Protocol) [35] transit: typically, a smaller ISP pays a larger provider for the traffic transited in both directions of the link between the two ISPs. Although the transit prices *per Mbps* decline steadily [29], the overall IP transit costs remain high or even increase according to industry analysts [8, 28, 36].

The problem of reducing the IP transit costs has attracted notable solutions of IXPs (Internet eXchange Points) [5, 21], IP multicast [6, 10, 19], CDNs (Content Distribution Networks) [9, 49, 56], P2P (Peer-to-Peer) localization [14, 62], and traffic smoothing [38, 43]. One property that these proposals share is their objective to reduce the amount of traffic that traverses transit links. Intuitively, the less traffic of an ISP flows through those links, the lower the cost is for the ISP.

This paper proposes *CIPT (Cooperative IP Transit)*, a different approach to reducing the cost of IP transit. Instead of altering the traffic that flows through the transit links, CIPT reduces the price of transit *per Mbps*: by jointly purchasing the IP transit, two or more ISPs reduce the transit prices per Mbps for *each* ISP involved in the CIPT.

While CIPT is a novel proposal in the context of the Internet ecosystem, group buying (tuangou) has been highly successful in other domains [39]. Similarly to the tuangou elsewhere, CIPT succeeds primarily due to subadditivity of prices [29, 60]. However, the benefits of CIPT depend also on burstable billing [23], different methods to account for bidirectional traffic, and other complex factors.

Relying on real inter-domain traffic and transit pricing, this paper estimates the gains from CIPT. We also propose Shapley value as a basis for sharing the gains among the CIPT partners so that to provide each partner with a strong economic incentive for the cooperation. Our evaluation of the aggregate and individual gains involves collection of the visual traffic statistics from 6 public IXPs with 264 participating ISPs, transformation of the visual images into a numeric format, and public-data validation of the property that peering and transit traffic have similar temporal profiles. Our analysis suggests that the expected relative savings of CIPT are in the range of 8-56% for the IXP-wide coalitions; in absolute terms, each of the partners may expect annualized savings from one thousand US$ for very small ISPs to several hundred thousand US$ for the few large ISPs. We also show that much smaller coalitions, with a half a dozen of members, can offer close-to-maximum savings. The main contributions of our paper can be summarized as follows:

- We propose CIPT, a simple strategy to reduce costs by purchasing IP transit jointly.
- We show that CIPT can be modeled as a cooperative game and that Shapley value provides an intuitive mechanism for cost sharing in CIPT.
- We use public IXP data to infer the traffic time series for several hundred (mostly regional and national) ISPs and use this information to assess the potential cost benefits of CIPT.

While our results on the CIPT cost reduction validate the potential of CIPT to be become a new viable element of the Internet ecosystem, the practical viability of CIPT also depends on other strategic and organizational issues. For example, if two ISPs are already engaged in a transit relationship, they are unlikely to agree on buying IP transit jointly from a third party. Also, the transit provider can strategically respond to CIPT by charging the coalition at higher prices per Mbps than the prices offered to an individual ISP. On the other hand, big transit providers can strategically adopt CIPT to squeeze out smaller transit providers: while a big transit provider might be unwilling to deal with a multitude of tiny customers, CIPT serves as a traffic aggregator and can reach the size attractive for the big transit provider; by selling IP transit to the CIPT, the big transit provider expands its customer base at the expense of the smaller transit providers which lose their individual tiny customers. Organizationally, CIPT faces a challenge of measuring individual traffic profiles accurately without violating traffic confidentiality. Whereas the above considerations can affect the size and composition of CIPT coalitions in reality, CIPT will probably not become the dominant mechanism for IP transit cost reduction. Still, we expect CIPT to gain broad presence in the Internet ecosystem, from small websites in a hosting facility to the level of nation-wide ISPs. However, data-driven assessment of all these additional issues lies beyond the scope of this paper. Similarly, while we propose Shapley value as a means for cost sharing in CIPT, evaluation of alternative solutions to CIPT cost sharing is a topic for future work.

The rest of the paper is structured as follows. Section 2 reviews the particulars of IP transit pricing and illustrates the CIPT potential with a simple numeric example. Section 3 formulates CIPT as a cooperative game. Section 4 explores CIPT cost sharing. Section 5 evaluates CIPT based on the public data. Section 6 discusses strategic, organizational, and other aspects of CIPT. Section 8 presents related work. Finally, Section 9 sums up the paper and its contributions.

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

Table 1: IP transit pricing rates of Voxel.

2. BACKGROUND AND MOTIVATION

The geographic location affects significantly the cost of IP transit. The IP transit prices per $Mbps$ per month range usually from \$5 to \$100 (we use \$ or US\$ to refer to U.S. dollars throughout the paper): the wholesale IP transit is typically priced under \$10 per $Mbps$ in most European and North American hubs but can exceed \$100 per $Mbps$ in Australia, Latin America and other remote regions of the Internet [3, 29].

Regardless of the geographic location, IP transit is subject to economies of scale and is priced subadditively: the prices per $Mbps$ are smaller for larger quantities of IP transit [29, 60]. Table 1 presents the current (as of January 2011) transit pricing rates of Voxel, a transit provider in North America [60]. The table reports the prices for different levels of CDR (Committed Data Rate), the minimum amount charged by the provider. For example, an ISP with IP transit needs of 300 $Mbps$ commits at the 100-$Mbps$ CDR level and pays pro rata \$3000 to Voxel but an ISP with IP transit needs of 700 $Mbps$ finds it more cost-effective to commit at the 1000-$Mbps$ CDR level and pays \$5000.

Burstable billing is another important aspect of IP transit pricing [23, 43]. To calculate the IP transit cost, the most commonly used method is to calculate the *peak* usage (typically through the 95th-percentile rule [23, 43]) and then the price function f is applied to

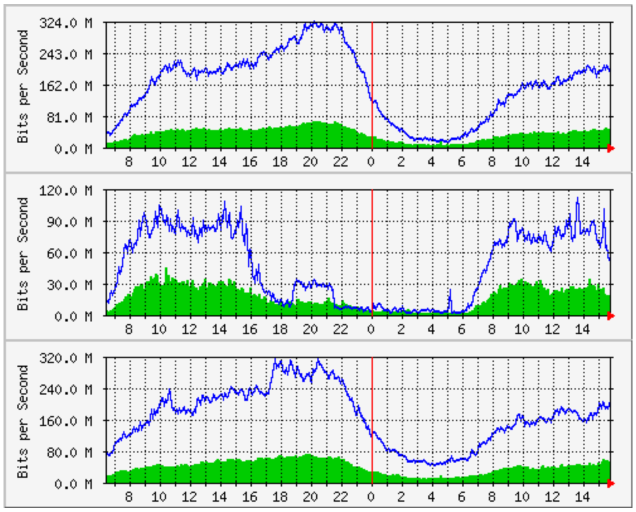

Figure 1: Demand statistics for partners P_1 (top), P_2 (middle), and P_3 (bottom) in the motivating example: the x-axes are in hours; the y-axes are in $Mbps$; the filled (green) areas depict the upstream traffic; the (blue) lines represent the downstream traffic.

the observed *peak* to calculate the resulting payment. The *peak* value is usually calculated separately for the upstream and downstream directions, and either sum or maximum of the two is used for billing. We refer to these two pricing models as sum and max models. Intuitively, the max model offers a larger opportunity for savings in cooperation because two ISPs with their traffic peaks in opposite directions can mutually benefit from the less utilized directions of each other.

Finally, it is worth noting that the the prices (per $Mbps$) of transit are in constant decline over the previous decade, with an average decay of around $25-30\%$ per year [29]. While this trend appears to be inevitable with the increase of the market competition, the total amount of interdomain traffic grows with a rate that outpaces the decay in prices. The recent paper from Arbor networks [37] reports an annualized interdomain traffic growth (on a set of 110 geo-diverse ISPs) of 44.5%. CISCO [15] and MINTS [44] report slightly higher annual growth figures, in the range of $50-60\%$. While a fair fraction of this growth is due to increased peering [5], there is still a consensus that most of the access/content providers do not see the reduction in their transit bill. To quote Erik Kreifeldt, a senior analyst for TeleGeography: "*... the growth offsets the price decline, so revenue* (of transit providers) *is more or less consistent or growing*" [36]. Similar observations have been made by several other business analysts [8, 28].

To illustrate the potential of CIPT, we consider a simple scenario of three partners[1] P_1, P_2, and P_3 interested in purchasing IP transit from the same provider. We assume the transit pricing rates as in Table 1, 95th-percentile burstable billing, sum model of accounting for bidirectional traffic, and traffic profiles plotted in Figure 1.

If the three partners purchase the IP transit separately, the individual traffic peaks (computed as the sum of the peaks in both directions) of P_1, P_2, and P_3 are at 379 $Mbps$, 130 $Mbps$, and 362 $Mbps$ respectively, and each of the partners commits at the 100-$Mbps$ CDR level. Thus, partners P_1, P_2, and P_3 pay respectively \$3790, \$1300, and \$3620 with the aggregate transit cost of \$8710.

On the other hand, if P_1, P_2 and P_3 use CIPT to buy the IP transit together, their aggregate peak traffic is 712 $Mbps$. By committing at the 1000-$Mbps$ CDR level, the CIPT pays \$5000. Thus, the cooperation reduces the aggregate transit cost of the partners by \$3710, or 43%. This significant cost reduction comes from two different sources:

1. Burstable billing – the 712-$Mbps$ peak of the aggregate traffic is lower than the 871-$Mbps$ sum of the individual traffic peaks; hence, the aggregate transit cost would reduce even if the pricing function were additive;

2. Subadditive pricing – the upgrade from the 100-$Mbps$ CDR level to the 1000-$Mbps$ one provides a lower price per $Mbps$ and thereby reduces the aggregate transit cost even further.

3. COOPERATIVE IP TRANSIT

In Section 1 we sketched the main idea of the CIPT. This section provides more details and discusses several aspects of the strategy.

We use term *Cooperative IP Transit* (CIPT) to refer to any cooperative mechanism in which two or more subjects purchase the IP transit jointly as a means for cost reduction. The subject interested in CIPT can be any Internet entity that buys IP transit; such entities include websites and hosting providers, as well as access, nonprofit, and content ISPs

The main incentive for forming a CIPT coalition is financial: each partner reduces its individual IP transit bill. The typical IP transit pricing makes it virtually impossible for a set of potential partners to increase their aggregate transit cost by buying the IP transit jointly. However, CIPT needs a reasonable mechanism to distribute the aggregate cost savings among all the CIPT partners. Furthermore, the aggregate and individual IP transit costs of the CIPT partners strongly depend

[1]We interchangeably use terms partner and player to refer to any ISP, hosting provider or any other entity interested in purchasing IP transit.

on a number of factors such as the IP transit pricing function, number of partners, their size, and temporal patterns of their traffic demands.

Formally, CIPT is a set of N partners. Each partner i of the CIPT has upstream and downstream IP transit traffic demands represented respectively by time series $u_i(t)$ and $d_i(t)$ where $i \in \{1, 2, \ldots, N\}$, and time t is measured in fixed-size time intervals with a typical interval duration of 5 minutes. The cost subject i pays for the transit, without participation in CIPT, is the function of these demand series:

$$C_i = F(u_i(\cdot), d_i(\cdot)).$$

After bundling of N subjects, the aggregate upstream/downstream demands are the sum of the corresponding individual demands:

$$u(t) = \sum_{i=1}^{N} u_i(t) \text{ and } d(t) = \sum_{i=1}^{N} d_i(t),$$

and the aggregate cost of the IP transit is

$$C = F(u(\cdot), d(\cdot)).$$

The 95th-percentiles of the upstream ($peak^{(up)}$) and downstream ($peak^{(down)}$) traffic are calculated, and the *peak* value used for billing is either the sum or max of these two values, depending on which of these two models (described in Section 2) is used. The transit cost of the coalition of these N players is then

$$C = F(u(\cdot), d(\cdot)) = f(peak)$$

where f is the pricing function decided by the IP transit provider. This pricing function is typically subadditive; see Section 2 for an example of such pricing function used by North American transit provider Voxel. As we will see, virtually always the overall IP transit cost of CIPT is strictly smaller than the sum of individual IP transit costs of all involved players:

$$\rho = \frac{C}{\sum_{i=1}^{N} C_i} < 1.$$

The relative savings $(1 - \rho)$ of the CIPT are influenced by several factors, with the two dominant being: (1) the subadditivity of the price function and (2) burstable billing through the 95th-percentile method. Namely, the subadditive pricing allows obtaining lower prices (per $Mbps$) when buying at larger quantities, which in turn allows savings for the involved players. Additionally, with the burstable billing, when two or more players have non-overlapping peak hours, their coalition would have the peak value strictly smaller than the sum of the peak values of the involved players. While players that serve similar user bases have similar temporal usage patterns (e.g. residential networks peak in evening hours, government/academic networks peak in early afternoon), the networks of different types experience their peaks in times that are far apart, which in turns allows for additional savings on top of bundling and buying-in-bulk.

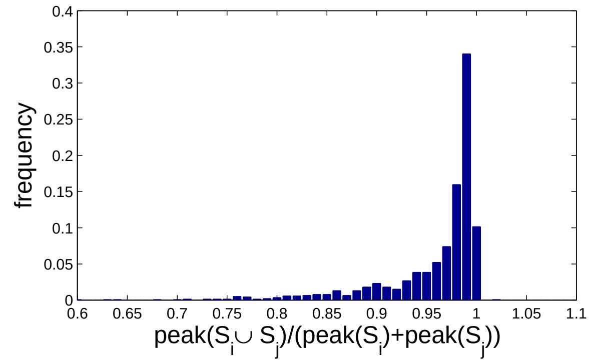

Figure 2: The distribution of ratios of the 95th-percentile of the union to the sum of the 95th-percentiles across all the pairs of ISPs from the Slovakian IXP.

Comment 1. *While this paper focuses on IP transit, the CIPT concept is relevant and straightforwardly applicable to cost reduction in other Internet business domains, such as IP transport and IXPs. As with IP transit, purchase of an IP transport link between two remote locations is also costly and subject to subadditive pricing. Multiple ISPs that need to reach the same remote location (e.g., an IXP) can reduce their IP transport costs by jointly buying a single IP transport link. Nonprofit IXPs constitute another instance of the CIPT concept: instead of buying IXP services from a third-party commercial provider, multiple ISPs can form a nonprofit IXP, cooperatively pay for the IXP infrastructure, and thereby reduce their peering costs.*

3.1 CIPT as a cooperative game

In this section, we briefly describe the concept of cooperative (or coalitional) games and show that CIPT can be seen as a cooperative game.

A cooperative game is characterized by set $\mathcal{N}$ of involved players and a cost function that maps the partitive[2] set of $\mathcal{N}$ to a cost value: $c : 2^{\mathcal{N}} \to R$. In the context of CIPT, set $\mathcal{N}$ is the set of subjects interested in purchasing IP transit. The cost function maps an arbitrary subset $S \subset \mathcal{N}$ to the cost of the IP transit that the coalition of players from S would pay. An important property of the IP transit model is that the price per $Mbps$ is a non-increasing function of the *peak*, due to the subadditive nature of the pricing model.

Additionally, for virtually any real-world subjects interested in purchasing IP transit, the peak traffic of the

[2]For set $\mathcal{N}$, the partitive set of $\mathcal{N}$ is the set of all subsets of $\mathcal{N}$ and is usually denoted as $2^{\mathcal{N}}$.

union of two subjects is smaller than the sum of the peaks of these two subjects. In case of measuring the peak as the maximal traffic, this is an obvious consequence of the fact that the maximum of the sum of two nonnegative functions (over the same domain) is not greater than the sum of the maximums of these two functions. If the peak is measured through the 95th-percentile method, there may be some irregular cases[3] in which the sum of the 95th-percentiles is smaller than the 95th-percentile of the union of the traffic of the two subjects. However, these situations are extremely unlikely to happen in regular setups as we demonstrate in Figure 2. There we plot the ratio of the 95th-percentile of the union to the sum of the 95th-percentiles across *all* the pairs of ISPs from the Slovakian Internet Exchange (SIX). The SIX and several other IXPs publish traffic statistics that each of their members (mostly regional ISPs) exchanges at the IXP, and this information represents valuable and useful proxy for estimating the traffic patterns (volume, peak-hour, peak-to-valley ratio, up/downstream traffic ratio, etc.) for the involved ISPs; see Section 5.1.2 and Appendix for more details.

Observation 1. The traffic patterns of subjects interested in CIPT are such that for (almost) all pairs of coalitions S_1 and S_2 of these subjects, the peak value of the union of the two coalitions is smaller than the sum of the peak values of these two coalitions.

As we elaborate above, Observation 1 is very intuitive and can be empirically validated for available data of traffic patterns. From now on, we assume that subjects involved in CIPT are such that this observation is true. In that case, cost function $c(\cdot)$ is indeed subadditive:

$$c(S_1) + c(S_2) \geq c(S_1 \cup S_2), \text{ for any } S_1, S_2 \subset \mathcal{N}. \quad (1)$$

4. COST SHARING IN CIPT

A key question in any cooperation scheme created for cost reduction reasons is how to split the aggregate costs of cooperation. As we saw in Section 3.1 the CIPT can be abstracted as a cooperative game which puts us in a position to use the rich set of analytic tools for solving the problem of cost sharing. There are many solution concepts for cost sharing in cooperative games, including the core, the kernel, the nucleolus, and the Shapley value [61]. While other solution concepts have attractive features, in the context of CIPT we find particularly appealing to use the Shapley value since it has several distinct important properties, i.e. the Shapley value: (1) *exists* for any cooperative game and is *uniquely* determined, (2) satisfies basic *fairness* postulates [53, 61], and (3) is *individually rational* i.e. each player in CIPT receives a lower Shapley value cost than what it would be if it did not participate in CIPT. One potential deficiency of the Shapley value is that in general it is computationally hard to compute it exactly. However, state-of-the-art techniques provide simple and accurate methods for Shapley value approximation, as discussed in Section 4.2.

4.1 Shapley value: definition

For a cooperative game defined over set $\mathcal{N}$ of N players and each subset (coalition) $S \subset \mathcal{N}$, let $c(S)$ be the cost of coalition S. Thus, if coalition S of players agrees to cooperate, then $c(S)$ determines the total cost for this cooperation.

For given cooperative game $(\mathcal{N}, c(\cdot))$, the Shapley value is a (unique) vector $(\phi_1(c), \ldots, \phi_N(c))$ defined below, for sharing the cost $c(\mathcal{N})$ that exhibits the coalition of all players. It is a "fair" cost allocation in that it satisfies four intuitive properties: efficiency, symmetry, additivity and null-player; see [53, 61] for exact definitions of these properties and more details. The Shapley value of player i is precisely equal to i's expected marginal contribution if the players join the coalition one at a time, in a uniformly random order. Formally it is determined by:

$$\phi_i(c) = \frac{1}{N!} \sum_{\pi \in S_N} \left(c(S(\pi, i)) - c(S(\pi, i) \setminus i) \right) \quad (2)$$

where the sum is taken across all permutations (or arrival orders), π, of set $\mathcal{N}$ and $S(\pi, i)$ is the set of players arrived in the system not later than i. In other words, player i is responsible for its marginal contribution $c(S(\pi, i)) - c(S(\pi, i) \setminus i)$ averaged across all $N!$ arrival orders π. Note that the Shapley value defined by Eq. (2) indeed satisfies the *efficiency* property:

$$\sum_{i \in \mathcal{N}} \phi_i(c) = c(\mathcal{N}).$$

4.2 Estimation of Shapley value in CIPT

While the Shapley value can be computed in a rather straightforward manner using (2), it is not practically feasible to employ (2) for $N > 30$. A number of methods have been suggested for accurate estimation of Shapley value, and in this paper we use a very simple Monte Carlo method, analyzed in [40], as follows.

Instead of calculating the exact Shapley value as the average cost contribution across all $N!$ arrival orders, we estimate the Shapley value as the average cost contribution over set Π_k of K randomly sampled arrival orders:

$$\hat{\phi}_i(c) = \frac{1}{K} \sum_{\pi \in \Pi_K} \left(c(S(\pi, i)) - c(S(\pi, i) \setminus i) \right) \quad (3)$$

The parameter K determines the error between the real Shapley value and its estimate: the higher K the lower

[3]For example, two subjects consuming 100 $Mbps$ 4% of the time each, one in the morning the other over night, and using 1 $Mbps$ the remaining 96% of the time will have their 95th-percentile equal to 1 $Mbps$, while their union would have 95th-percentile equal to 100 $Mbps$.

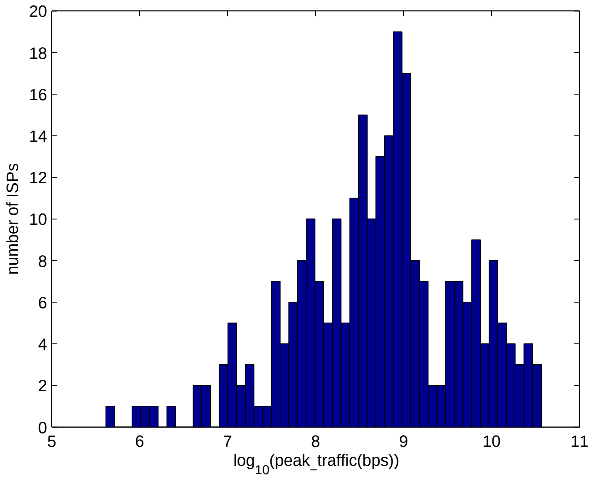

Figure 3: The distribution of the peak traffic rates across all 264 ISPs: median: 560 $Mbps$**; mean:** 2.9 $Gbps$**.**

the error. So basically, one can control the accuracy of the estimator by increasing the number of sample permutation orders. We observe in our datasets of traffic demands that the value of $K = 1000$ provides errors of under 1% across all the CIPT players, and in the rest of the paper we use $K = 1000$ for the computation of the Shapley value.

5. EVALUATION

In this section we quantify various factors that impact CIPT by using traffic information from 264 (mainly national and regional) ISPs. In Section 5.1 we describe the dataset and pricing model(s) used. In Section 5.2 we evaluate the potential savings of CIPT on country-wide (IXP-wide) collaborations and show that significant savings could be expected both in relative and absolute terms. In Section 5.3 we augment this analysis by empirically showing that even small, single-digit coalitions, can yield close-to-optimal savings, by demonstrating a law of diminishing returns for the savings as a function of the coalition size. Section 5.4 analyzes the per-player savings and shows somewhat expectable trends that the larger the player is, the larger are its absolute savings, but the smaller its relative savings are. Finally in Section 5.5 we analyze the effects of collaboration between geo-diverse players and present an analytical upper bound on the savings as a function of the time difference in their peak-hour periods.

5.1 Dataset description

Although data for the traffic patterns of many ISPs is often kept confidential some public Internet eXchange Points (IXP) report upstream and downstream demand time series for the traffic exchanged by every member of the IXPs. Those that do it are listed in the Table 2. This traffic statistics data is typically given in the form of mrtg images [11], similar to those shown in Figure 1. Overall we collected the information for 264 ISPs, with the traffic peak distribution as shown in Figure 3. While the information about the traffic exchanged at the public exchanges is obviously a valuable piece of information, it is not straightforward how to use this information to estimate the transit usage of the ISPs. In Section 5.1.2 we use a small set of ISPs that make their detailed traffic information public, to show that the IXP related traffic is a good proxy for estimating the transit part of the interdomain traffic, at least for some ISPs. Before that, we elaborate on the data collection in the following Section 5.1.1.

5.1.1 *Dataset collection*

We started by manually inspecting the webpages of the medium-sized and large IXPs [26]. A majority of these IXPs publish their aggregate traffic statistics, summed across all the members, but some also make public the detailed traffic statistics of their members. We identified several IXPs that do so; they are listed in Table 2. We then crawled the websites of these IXPs and collected per-member traffic information. This per-member traffic data is typically given in the form of visual images, similar to those in Figure 1, produced as the outputs of the standard tools for traffic visualisation: mrtg/rrdtool [11]. To convert the information into a numeric form, we built a piece of software that takes as input a mrtg/rrdtool image and outputs the numeric array representing the upstream/downstream traffic time series. This operation of transforming the .png images to numeric data required serious effort in the domain of optical character and function recognition. We plan to release for public use both the numeric data itself and the code for transforming mrtg/rrdtool images into the numeric format.

5.1.2 *From IXP data to IP transit traffic*

Most ISPs consider the data of their networks as very confidential and are reluctant to share it with third parties. However, some ISPs share publicly large amounts of operational information data. In particular, several European ISPs serving academic institutions have shared publicly on their websites detailed picture of both their network infrastructure and utilization of their networks. Those that we identified are HEANET (Ireland) [32], SANET (Slovak Republic) [52], CESNET (Czech Republic) [12], GRNET (Greece) [31]. We inspected the peering and transit traffic for those four ISPs and found, somewhat expectably, that the peering traffic pattern is a good first-order indicator of the transit traffic. We found that, in those 4 ISPs, peering

IXP	acronim	# of members	peak ($Gbps$)	average ($Gbps$)	95th-pct effect (sum/max)	skewness
Neutral IX (Prague)	NIX	54	116	76	4.3%/29.1%	0.76
Slovakian IX	SIX	52	42	23	15.4%/44.9%	0.27
Israel IX	IIX	17	2.1	1.38	14.3%/40.6%	0
Finish IX	FICIX	25	32	19	6.7%/23.1%	0.48
Interlan (Bucharest)	Interlan	63	22	11	14.3%/37.8%	0.12
Budapest IX	BIX	53	152	92	3.6%/27.8%	0.84

Table 2: Basic stats on the used IXPs.

corresponds to 35-40% of the total traffic, with the remaining 60-65% being transit. Additionally, we observe that peering and transit traffic follow very similar temporal patterns: their growth and decay periods coincide, they peak in the same time, have similar peak-to-valley ratios, etc.; see Appendix A for more details. In some sense, such behavior is not very surprising: given that the demand is predominantly created by humans, both transit and peering traffic demand are driven by the same end-user activities.

Consequently in our analysis, we approximate the transit traffic of ISPs (belonging to corresponding IXPs) with their peering traffic (information that is publicly available) multiplied by a factor γ that determines the relative weight of the transit vs. peering traffic. We believe that, in spite of this relatively crude approximation, this first-order estimation provides a good starting point for evaluation of CIPT and factors that affect it: relative sizes of the players, temporal effects, peak-to-valley ratio, etc. In section 5.2 we describe expectable savings of CIPT for a range $\gamma \in [0.5, 4]$. However, in the Sections 5.3-5.5 (which analyze the cost-sharing, coalition size, and geo-diversity), we fix $\gamma = 1.5$, that corresponds to transit vs. peering traffic ratio of $60 : 40$ as suggested by our analysis in Appendix, for medium-sized European countries with a single dominant IXP (the case of our 6 IXPs).

5.1.3 Pricing model

In the following evaluation we use the Voxel pricing model (described in Section 2) with prices given in Table 1 and upstream/downstream traffic billed with either sum or max model. In section 5.2 we describe the results of a comparative study of both sum and max models. In the Sections 5.3-5.5 we will focus on the sum pricing model (the more conservative one in terms of cost reduction) for the analysis of cost-sharing, coalition size and geo-diversity.

5.2 Aggregate savings

In this section, we evaluate the aggregate potential savings of the IP transit costs for the coalitions consisting of *all* members of IXPs listed in Table 2. Following the discussion in Section 5.1.2, we approximate the IP transit traffic patterns by the traffic exchanged at these IXPs multiplied by constant $\gamma \in (0.5, 4)$; this constant represents the ratio between the transit and IXP traffic volume. While this approximation is rather crude, it nevertheless captures the main features of the ISP: relative size, peak-hour period, upstream-to-downstream ratio, etc. For example, $\gamma = 0.5$ corresponds to the case where the peering traffic amounts to $1/(1+\gamma) = 2/3$ of all the traffic of the ISP (as in Japan[13] and other very localized markets), while $\gamma = 4$ corresponds to the case where $1/(1+\gamma) = 20\%$ of the total ISP traffic is exchanged at the IXP, and the remaining 80% is transferred through transit (this situation is common in small markets [3]). The empirical evidence of few European ISPs discussed in the Appendix, suggests that in medium-sized European countries with one dominant IXP: $\gamma \in [1.5, 2]$.

We stress again that the purpose of this evaluation is to shed some light for the potential savings of CIPT rather than computing accurate bounds of the savings. Such exact saving estimates strongly depend on various factors and should be calculated on a case-by-case basis.

For each of the 6 studied IXPs, Figure 4 reports the expected savings on the IP transit bill, both relative and absolute, in both the sum and max models. We see that the relative savings are in the range of 5-70% depending on the relative size of the IXPs, and several other factors. This relative savings are strongly impacted by the size distribution of the involved ISPs. Namely for those IXPs that have several large ISPs that dominate the traffic (and the costs), the relative savings of CIPT are low, because these large ISPs already receive the lowest price per $Mbps$. To illustrate that this is indeed the case we define the *skewness* factor as the fraction of the traffic generated by the players with peak traffic greater than 10 $Gbps$. We see from the Table 2 that for the IXPs with a low *skewness* of under 0.3 (SIX, IIX, and Interlan), the expected relative savings are considerably higher than those of the others.

Remember that the savings of CIPT come from two properties of the IP transit model: price elasticity and the 95th-percentile billing. A crucial question in this context is to quantify the effects that these two properties have on the CIPT savings. For this purpose we identified what the relative savings would be without the elasticity of the prices, i.e., if the price per $Mbps$

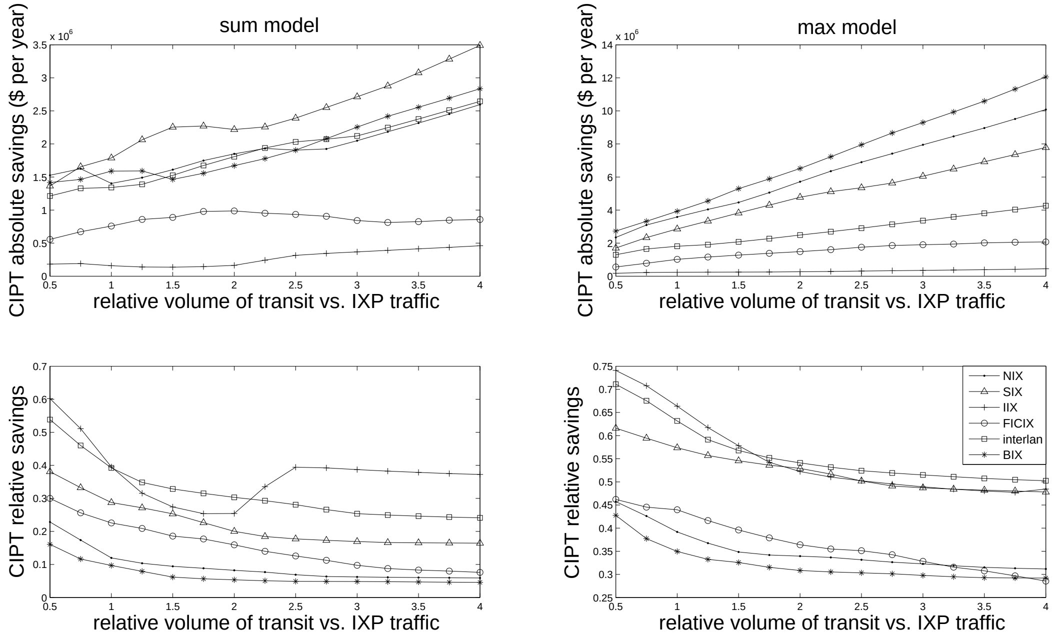

Figure 4: The absolute and relative savings as a function of the ratio between the transit and IXP traffic volumes.

would be constant independent of the usage level. Such savings would come exclusively from the reduction in the 95th-percentile. The results are given in Table 2 under the column *95th-pct effect*. From this table we can conclude that both properties (price elasticity and the 95th-percentile billing) influence the total savings. However the exact breakdown of the impact of these two properties on the total savings depends on other factors.

The decreasing trend of relative savings can be observed in both sum and max pricing models. This is the consequence of the fact that the players with large volumes have smaller opportunities for large relative savings by CIPT (as they already experience low per $Mbps$ price). The relative savings are, however, bounded from below by the quantity *95th-pct effect* reported in Table 2 for both sum and max pricing models.

We conclude this analysis with an observation that these 6 (medium-sized European) countries hosting these IXPs, have such traffic locality that around 40% of the traffic stays inside the country, and is exchanged by peering (mainly through the dominant IXP), while the remaining 60% of the traffic uses IP transit (see Appendix). This corresponds to the value $\gamma \approx 1.5$. Using this value of γ, we conclude that the expected relative savings in IP transit costs for the IXP-wide CIPT coalitions are in the range of 8-35% (in the sum model) and 32-56% (in the max model).

5.3 Coalition size

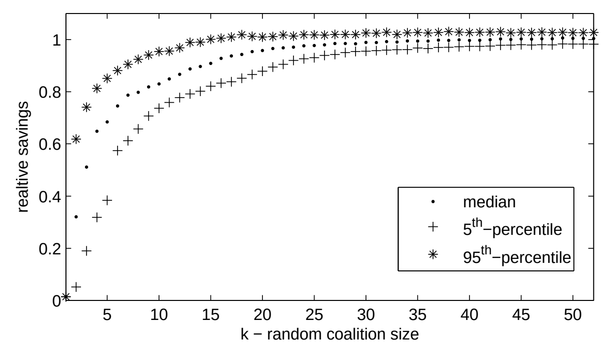

Figure 5: Relative (as fraction of the savings obtained in the grand coalition) per-player savings for smaller coalitions.

In previous section, we analyzed the potential savings of coalitions that include *all* members of the corresponding IXPs. While such coalitions offer significant savings in terms of IP transit costs, coordination of such large coalitions may be cumbersome. In this section, we show that much smaller coalitions can offer savings comparable to those of the large coalitions. For that reason, we take the Slovakian IXP (SIX) with $N = 52$ members, and for each $k \in \{1, 2, \ldots, N\}$ we analyze the per-player savings from participating in the coalition of k random members of SIX. The pricing model is sum, and $\gamma = 1.5$. The results for other IXPs, max pricing model and other choices of γ are very similar, hence we omit them for brevity.

In Figure 5 we report the median, 5th-percentile and 95th-percentile savings, relative to the savings obtainable from the grand coalition of all $N = 52$ members. Since analyzing the statistics across all 2^{52} subsets is infeasible, we report the results obtained by sampling: for each member i and each coalition size k, we pick random 100 subsets of size k, that contain member i. From Figure 5 we can observe the law of diminishing returns: relatively small coalitions provide savings very close to the savings of the large coalitions, and, by adding more members to the coalition, the incremental savings are decreasing. In particular, even with as few as $k = 3$ members, one can expect savings that are half as large as the savings obtainable by the coalition of all $N = 52$ members. With $k \geq 10$ members, the median CIPT savings are greater than 80% of the savings obtainable the grand coalition.

Note that the savings grow as the coalitions become larger. This is the consequence of the basic property of the CIPT cooperative game: the cost function is subadditive, as seen in Ineq. (1). In other words, by adding a member, the coalition is better off. Also, note that for some ISPs, participating in some smaller coalitions may be more beneficial than participating in the grand coalition (the relative savings > 1).

We stress that the results of this section are for random coalitions. By careful cherry-picking the most appropriate partners, one can obtain even higher savings, as the 95th-percentile of the savings in Figure 5 can suggest. However, such optimization is out of scope of the present paper.

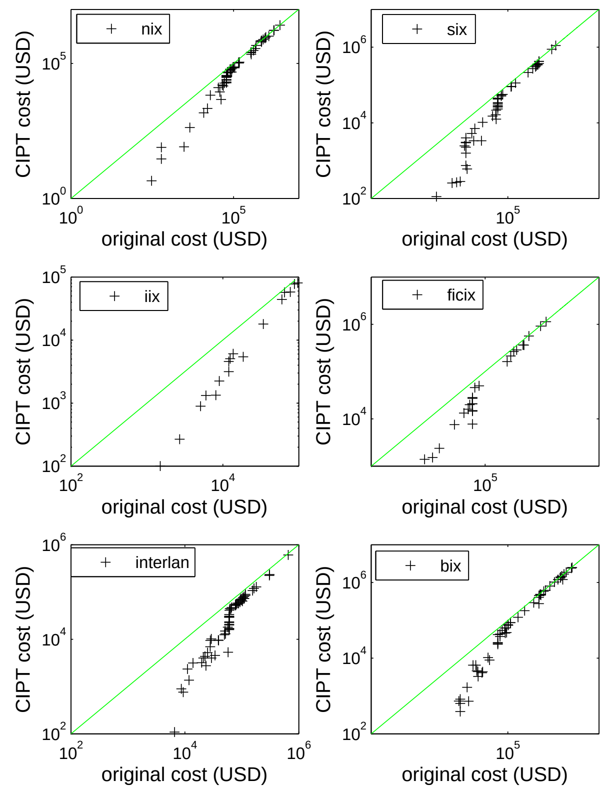

Figure 6: The original annual costs versus CIPT costs (Shapley value) across all the ISPs from the 6 IXPs.

5.4 Per player savings

In this section we look at the per-member savings for each of the involved ISPs when it participates in the IXP-wide CIPT. Following the reasoning described in Section 5.1.2, the γ factor used for scaling of the transit traffic is set to 1.5, and the pricing model is the more conservative `sum` model. As we elaborate in Section 4, each member of the coalition is assigned a cost equal to its Shapley value. The CIPT costs (across all ISPs) are depicted in Figure 6 against the original IP transit annual costs. Figure 7 shows the absolute annual savings (the difference between the original IP transit costs and CIPT costs) for all ISPs in these 6 IXPs.

We can observe two trends in Figures 6 and 7. First, the absolute savings typically grow with the size of the ISP. This is a consequence of the fact that having a large ISP in a coalition typically implies lower per $Mbps$ costs which in turn increases the contribution of the ISP to the coalition, and is reflected in the computation of Shapley value, eq. (2). In contrast to this increasing trend of the absolute savings, there is another interesting property of the CIPT cost allocation. Namely the relative savings of CIPT (the ratio of the absolute savings of CIPT to the original IP transit costs) typically see a decreasing trend as a function of the ISP size. This feature (decreasing trend of the relative savings) is strongly connected with the nature of Shapley value as a cost allocation strategy but is also related with the fact that the peak-hour of the coalition is predominantly determined by the large ISPs. This means that large ISPs that join already large-enough coalitions (those that reached close-to-minimum price per $Mbps$) do not bring large benefit to the coalition and consequently implying low relative-gains for these ISPs.

5.5 Cooperation between remote subjects

So far, our analysis was concerned with the ISPs operating in the same geographic area, and consequently having close peak hours. In such scenarios, the savings are mainly impacted by the price elasticity rather than the subadditivity of the 95th-percentiles. In this section we investigate potential savings of collaboration

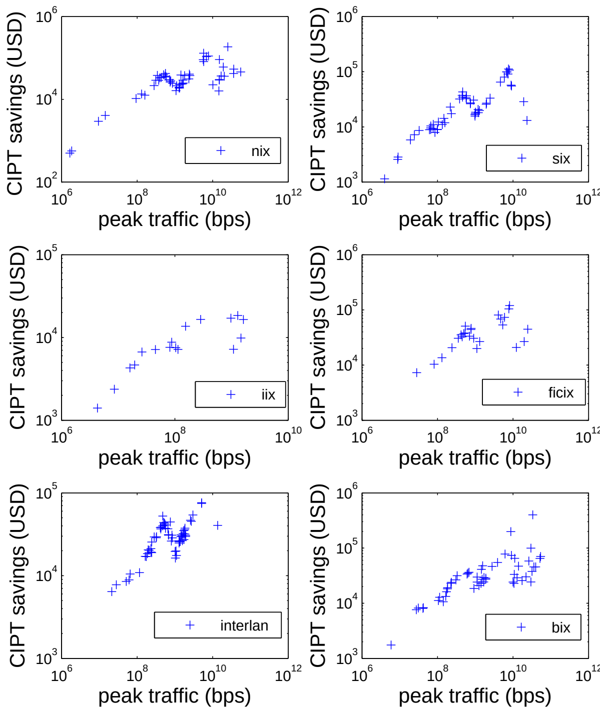

Figure 7: The absolute annual savings for all the ISPs from the 6 IXPs.

between geographically distant players. Collaboration between geographically distant players is possible only for large players. Only then the long-distance transport becomes cheap enough to make the CIPT economically viable [29]. Such long-distance transport to major (cheap) Internet hubs is not uncommon method for ISP cost optimization. For example, four largest IXPs: DE-CIX, AMS-IX, LINX, NYIIX, host ISPs from more than 40 different countries, each.

Additionally, cooperation between very remote subjects (say, more than 6 time zones), may strongly impact the performance in terms of increase of propagation delays. Some delay-sensitive applications (voice, gaming, etc.) may find such increase in delay unacceptable. Therefore, CIPT between very remote subjects is reasonable only for the traffic that is not delay sensitive (content, p2p, etc.) which indeed represents the majority of the Internet traffic [38, 43].

To analyze the potential savings in such a setup we look at the potential savings of collaborations with *two* players. Once *all* the players are large enough to receive the minimum per-$Mbps$ price, the coalitions with more than two players are not bringing large marginal benefits in terms of price reduction. Thus we here focus on 2-player coalitions. To assess the potential savings in such cases, we take all $M = 93$ ISPs from our 6 IXPs with *peak* traffic greater than 1 *Gbps*, and shift each one of them for a random (uniformly) number of time zones. For each of the $M(M-1)/2$ pairs, we evaluate the relative savings of the coalition: $1 - cost(CIPT(i,j))/(cost(i)+cost(j))$ and plotted it against the time difference in Figure 8. One can observe the trend: the further away the two players are, the greater the opportunity is for the CIPT savings. In Figure 8 we also depict the bound

$$g(\psi) = \frac{1 - |\cos\frac{\psi}{2}|}{2}, \qquad (4)$$

where $\psi = \frac{time\_difference}{24} 2\pi$ is the scaled time difference. We prove the upper bound on the relative savings in a simple model in which the demand curves are modeled as sin-waves (see below). One can observe that the relative reduction in the 95th-percentile for a coalition of two players is in the range of $[0, 0.5]$, in line with the model predictions. However, the expected savings appear to be larger as the time difference grows, and peak when two ISPs are 12 time zones apart. To explain and quantify this property we employ a simple trigonometric model where the demand pattern of the ISP is modeled as a sin-wave function. The following proposition characterizes the expected reduction in the peak traffic from CIPT collaboration between two players with non-coinciding peak hours:

PROPOSITION 1. *Let two players have demand given by*

$$D_i(t) = A_i \cos(2\pi \frac{t - M_i}{24}) + B_i, \quad t \in [0, 24) \text{ hours.}$$

where B_i is the mean traffic intensity, $A_i + B_i$ is the peak traffic intensity, and M_i is the peak hour of player i. By creating a CIPT coalition between these two players, the relative reduction in the peak is equal to:

$$G_{12} = 1 - \frac{B_1 + B_2 + \sqrt{A_1^2 + A_2^2 - 2cos\psi A_1 A2}}{B_1 + B_2 + A_1 + A_2} \leq g(\psi),$$

for $\psi = \frac{M_1 - M_2}{24} 2\pi$, the scaled time-zone difference, and $g(\psi)$ defined in (4)

PROOF. Omitted for brevity. $\square$

6. CIPT BEYOND THE COST SHARING

Section 5.4 presented a compelling evidence that CIPT with Shapley-value sharing of transit costs offers significant benefits to the CIPT partners. While the economic incentives are crucial for CIPT being viable, the viability is a topic with multiple dimensions. Without pretending to be comprehensive, this section discusses other aspects of CIPT such as its organizational embodiment, physical infrastructure, performance, traffic confidentiality, interdomain routing, relationship with transit providers, and social impact.

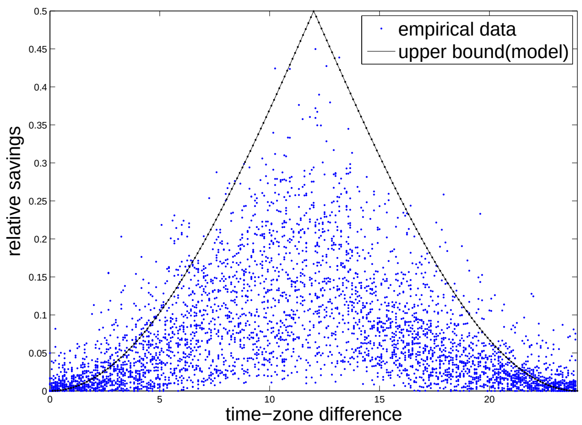

Figure 8: Relative savings between large remote subjects coming from the 95th-percentile subadditivity.

Organizational embodiment: CIPT is an innovative mechanism for reducing transit costs. Among other cost-reduction mechanisms, peering is similar to CIPT in its cooperative nature and commonly organized as a nonprofit IXP. In our vision for CIPT as an organization, a typical arrangement is also a nonprofit organization. The nonprofit status of a CIPT promotes a valuable marketplace image of its neutrality and fair treatment for all its partners. In such an organization, partnership fees are used only to recover the costs of operating the CIPT and expected to be insignificant in comparison to the transit cost reductions provided by the CIPT. While the nonprofit arrangement looks the most suitable, deviations are quite possible and even likely; as with some existing IXPs, some CIPTs might operate as government or commercial organizations.

Physical infrastructure: The physical implementation is another issue where CIPTs can benefit from the IXP experience. For buying IP transit in bulk, a CIPT needs to concentrate traffic of multiple ISPs in one location. The physical infrastructure of any IXP already supports such concentration for peering purposes. Moreover, some IXPs diversify their service portfolio by offering access to transit providers. For example, Vancouver Transit Exchange is an IXP that also hosts transit providers and thereby enables an ISP to satisfy its peering and transit needs at the same location [33]. A CIPT can be implemented as a further diversification of the IXP service portfolio. By leveraging the physical infrastructure of an existing IXP, the CIPT can keep its operational costs low.

Performance: A CIPT and its transit provider sign a contract for IP transit. The contract is expected to be of the same type as existing contracts between an individual ISP and its transit provider. In particular, the contract includes an SLA (Service Level Agreement) [57] stating the maximum outage duration, packet delay, jitter, and loss rate for the CIPT traffic. The SLA also specifies financial compensations by the provider if the latter fails to provide the CIPT with the agreed performance. In reality, SLA violations are likely to be rare. Whereas the performance levels of traditional inter-provider SLAs are very similar, having a single SLA for the multiple-partner CIPT is not problematic. Also, the typical SLA metrics of packet delay, jitter, and loss rate are such that the traffic of individual CIPT partners can inherit the performance levels of the CIPT aggregate traffic without any special technical support. Furthermore, the CIPT and its individual partner can sign a separate bilateral agreement on performance issues.

Traffic confidentiality: While it is feasible to formalize traffic metering and billing for a CIPT by means of bilateral agreements between the CIPT and each of its individual partners, the bill of a partner depends on the traffic of the other partners. Some academic ISPs – such as the aforementioned HEANET, SANET, GRNET and CESNET – reveal their transit and peering traffic. However, a typical commercial ISP tends to be more secretive and does not disclose its traffic patterns. To alleviate the privacy concerns, a CIPT can keep the traffic profiles of its partners confidential and incorporate an internal audit system for verifying the correctness of traffic metering and billing for each partner. Note that the confidentiality undermines the formation of most effective CIPTs. Making the traffic profiles of ISPs and CIPTs public would help in determining the best matches between CIPTs and their potential partners. In general, the overall efficiency of the Internet industry would benefit from more transparent traffic practices.

Interdomain routing: With BGP (Border Gateway Protocol) [50] being a de facto standard protocol for routing between ASes (Autonomous Systems), we see no technical complications with CIPTs from the interdomain routing perspective. A CIPT can acquire a separate AS number for inclusion into its BGP path announcements. Alternatively, as in the case of some IXPs, the partners of a CIPT can agree to use the individual AS number of one (typically, prominent) partner in all BGP announcements by the CIPT.

Relationship with transit providers: The costs saved by the CIPT partners are the revenues lost by the transit provider. Hence, transit providers are likely to perceive CIPTs negatively. On the other hand, IP transit is a competitive market with low levels of government regulations. By refusing to serve a CIPT, a transit provider would hurt mostly itself because the CIPT would then take its transit business to another provider. As with IXPs in some countries, transit providers can lobby their national governments to outlaw CIPTs. However, as in the IXP cases, the success of the legal actions is likely to be limited and temporary. Without effective means to suppress the CIPT innovation, transit providers will need to coexist with CIPTs and find their own new ideas to compensate for the diminished revenues.

Social impact: Looking beyond the economic interests of individual parties, the overall social impact of CIPTs appears positive. In particular, CIPTs are beneficial for narrowing the digital divide between the developed countries and poorer world which lies on the Internet edges and does not own a transit infrastructure for reaching the Internet core. In places like Africa, IP transit (and IP transport) is more expensive but the ability to pay for it is lower. Like with IXPs that have positively affected Africa by exchanging its traffic locally rather than through North America or Europe, CIPTs can benefit this and other developing regions by making the access to the Internet and its information more affordable [3].

7. CIPT: A STRATEGIC PERSPECTIVE

Previous sections have analyzed feasible gains for CIPT members without considering possible strategic behaviours. It would be however naïve not to expect strategic reactions to and within CIPT. While it is not the objective of this section to provide a comprehensive account of all the potential strategic implications it is worth making some considerations. Future work will be devoted to deepening this analysis.

Strategic scope is particularly relevant at two levels. Within the CIPT, members might be greedy, trying to obtain extra benefits by leveraging their bargaining position against other members of the CIPT. Also the participation in the CIPT might be affected by the presence of existing or potential customers/providers or peers. Such issues related with CIPT formation and participation will be examined in subsection 7.2. CIPT strategic implications are not restricted to potential or current members: other transit providers and transit buying ISPs could react to CIPT formation. In the later case new coalitions of customer ISPs could replicate the CIPT scheme, reinforcing CIPTs bargaining position. More interestingly, reactions by transit providers are expectable. Both the transit provider engaged in commercial relationship with a CIPT and its competitors are likely to react among each other and towards the CIPT. These reactions will be studied in the next subsection (7.1).

7.1 Transit providers« strategic perspective

While it seems obvious that group-buying can be desirable from the point of view of customer ISPs, transit provider's perspective is not that straightforward. Although apparently counterintuitive, transit providers' strategic behaviour underpins the feasibility of CIPT for at least two reasons. Regardless of whether the transit provider enjoys a monopolistic position or not, aggregation of the demand of smaller customers might allow a transit provider to directly provide IP transit to such customers bypassing intermediaries. Secondly, if the transit provider is not a monopolist, i.e. there is at least another transit provider, the former can attract customers through CIPT to the detriment of its competitors.

Due to the existence of large economies of scale, provision of transit to small networks tends to be unattractive for big transit providers. Instead middle-sized networks usually act as resellers providing arbitrage between small and big networks. Aggregation of small customers can make direct selling of transit to small networks economically attractive for bigger transit providers. Such strategy might yield benefits for both the transit provider and the customers aggregated in the CIPT. CIPT members benefit from joining the CIPT by committing to a higher CDR and hence enjoying lower prices as it has been shown throughout the paper. Conversely, increased revenues for the transit provider arise because without reducing the total amount of sales, average price can be increased. In one hand the original intermediary, the transit reseller bypassed by means of the CIPT, reduces its CDR facing hence a higher price. In the other hand the revenues captured by this network now accrue to the big transit provider.

However the ability of a transit provider to undertake these strategy is determined by its market power, i.e. the ability of the network to raise prices without undermining profits. If the aggregation of smaller customers is not attractive enough for the transit provider and it enjoys a monopolistic position, CIPT wont take place unless the total revenue for the transit provider is the same as if each ISP would buy transit independently. In such a case the total cost faced by the coalition is the same as if each of the members would buy transit independently. The gains of a member can only happen at the expense of another one. Therefore from the point of view of the transit provider CIPT would be either equivalent to a situation without group-buying either it wont happen.

Differently, when there is competition a transit provider can expand its revenues at the expense of the revenues of its competitors. Competing transit providers could react accepting CIPT as well. This interaction scheme could eventually lead to a pricing war between competitors driving prices to the minimum so firms just recover their costs. CIPT members would enjoy an increased bargaining power and corresponding benefits due to the drop in transit prices. Price competition among transit providers and customers deciding the CDR that maximizes their expected utility resembles to the typical model of Bertrand Competition, where firms produce a homogeneous good and compete by setting prices while consumers decide quantities. In such a situation the firm offering the lowest price attracts all the customers of the market, leading to a pricing war between competing firms that drives prices to the minimum so firms just covers their respective costs. This is the so called Bertrand Paradox. However in practice various circumstances usually foreclose this extreme from happening. ISPs favor multihoming, besides costs structures, capacity, offered quality and coverage of transit providers is not homogeneous. Consequently attracting all the traffic of the market by a mere decrease on price is not very realistic.

Reduction of transaction inefficiencies, demand uncertainty, heterogeneity among buyers and production postponement when combined with scale economies underpins other arguments usually adduced to demonstrate that group-buying can also lead to gains for the seller [?]. However in the present context reduction of transaction costs, net of coordination costs, are not likely to be big enough to incentivize the transit provider. Postponement of production is not viable because infrastructure is already deployed and the costs are hence sunk costs. Even though there is high uncertainty, demand heterogeneity and the provision of IP transit is subject to large economies of scale, transit providers can easily discriminate prices through quantity discounts as it is usually the case. Nevertheless the two aforementioned reasons, small customer's demand aggregation and competition among transit providers clearly justifies the feasability of CIPT.

7.2 Strategic issues within the CIPT

The composition of the members of the CIPT is also subject to strategic analysis. Two aspects are specially interesting: CIPT formation, i.e. which networks would join, and cost sharing, i.e. how the members of the CIPT could try to exploit the cost sharing mechanism, Shapley value, on their behalf.

From the point of view of coalition formation it is specially interesting looking at whether possible members are peers or customer and provider, and whether they belong to the same hierarchy level or not.

Reduction of transit costs is an incentive both for peering and for CIPT. By reducing transit costs CIPT might deter from joining the coalition networks with established peering agreements since it would reduce the value of peering, i.e. cost reductions due to peering would diminish. Similarly CIPT members would witness a reduction in the incentives towards peering and consequently demand for transit could grow, increasing transit provider's revenues. In as much as CIPT reduces peering and correspondingly increases transit, it poses another incentive for big transit providers to promote it.

By aggregating their demands CIPT members can successfully bypass intermediaries and obtain a direct connection with the bigger transit provider at a lower cost than through these middlemen. Hereby such intermediaries do not appeal as promoters of CIPT. Nevertheless once established the coalition and part of their traffic forgone it could be at their interest joining the coalition in order to minimize the damage.

This analysis has so far considered only a static setup. From a dynamic point of view, even if traffic contributions of coalition members were perfectly observable and cost sharing mechanisms worked without distortions time variations of traffic contributions could alter expected benefits from CIPT participation, increasing in turn uncertainty. If no commitment mechanisms exists ISPs could join declaring small traffic contributions and alter that in the future obtaining larger costs reductions at the expense of other members. A mechanism to reduce uncertainty in which CIPT members make a temporal compromise could be easily enabled to avoid this drawback.

8. RELATED WORK

In presenting and evaluating CIPT, we already mentioned the essential background information. This section takes a broader look at related work.

Our study of CIPT starts with the observations that interdomain traffic grows and that IP transit costs are high. The traffic growth is a long-term trend [13, 37], even though the main application fueling the growth has been changing from web browsing [25, 47] to P2P file sharing [51, 59] to video streaming [48, 63]. The recent investigation of 110 geographically diverse ISPs estimates the annual rate of the interdomain traffic growth at 44.5% [37]. Other reports cite even higher annual growth rates in the range of 50-60% [15, 44]. Whereas the IP transit is a competitive business, the transit prices per $Mbps$ decline [20] but at lower rates of about 25-30% per year [29]. In spite of the falling IP transit prices, ISP business analysts agree that the overall IP transit costs remain high or even increase [8, 28, 36].

The existing approaches for reducing the transit costs include ISP peering, IP multicast, CDNs, P2P localization, and traffic smoothing. Peering [5, 21] enables two ISPs to exchange their traffic directly, rather than through a transit provider at a higher cost. To disseminate data to multiple receivers, IP multicast [6, 10, 19] duplicates packets in IP routers and thereby reduces transit traffic. While IP multicast requires router support from transit providers, CDNs [9, 49, 56] and P2P systems duplicate data on the application level. Whereas a single company controls a CDN, a P2P system consists of independent hosts, and P2P localization [14, 62] strives to reduce transit traffic without undermining the system performance. Even if the transit traffic preserves its volume but is redistributed within the billing period to peak at a lower value, the transit costs decrease due to the burstable billing [23]. An ISP can do such traffic smoothing with rate limiting [43] or in-network storage for delay-tolerant traffic [38]. Unlike the above approaches that modify the transit traffic, CIPT reduces the transit costs without altering it.

CIPT can benefit from multihoming [1, 2] by connecting to multiple transit providers. While the connection reliability is a traditional rationale for multihoming, the latter also offers interesting trade-offs between performance and costs [30].

We view CIPT as a coalition and use the Shapley value [53] for sharing CIPT costs. Shair [34] is a cooperative system for a different application of sharing mobile phone minutes that enables phone users to share the committed but unused minutes. Cooperative approaches have also been studied for cost sharing in IP multicast [4, 27] and interdomain routing [42, 54, 55]. The game-theoretic analyses of the Shapley-value mechanism [4, 27, 46] highlight its group-strategyproofness and other salient properties but identify its high computational complexity. Despite the computational complexity, various proposals of traffic billing between ISPs [41], incentives in P2P systems [45], and charging individual users by access ISPs [58] rely on the Shapley value. Unlike the above applications of IP multicasting, ISP billing, P2P incentives, and individual user charging which involve a large number of parties, CIPTs are likely to be small in size. For CIPTs with few dozens of partners, the exact computation of the Shapley value is computationally feasible. Our evaluation of CIPTs uses the Monte Carlo method to estimate the Shapley value accurately [40].

As a new element of the Internet ecosystem, CIPT diversifies the means for the economic tussle between Internet stakeholders [16, 17]. Network neutrality refers to potential restrictions on ISP traffic management [18]. Similarly to peering or content caching [20], CIPT reduces transit costs without violating the network neutrality.

9. CONCLUSIONS

In spite of the steady decline of IP transit prices, the IP transit costs remain high due to traffic growth. Over the previous decades a number of solutions have been suggested to reduce these IP transit costs, including settlement-free or paid peering, IP multicast, CDNs, and P2P localization.

In this paper we propose an alternative cost-reduction technique of Cooperative IP Transit (CIPT), that in contrast to the existing solutions does not alter the traffic. Namely, CIPT utilizes tungu, or team-buying, for IP transit. The savings in CIPT come from two distinct yet ubiquitous properties of the IP transit pricing model: price elasticity and burstable billing. Our data-driven analysis suggests that significant savings can be expected from using CIPT. We are confident that the potential savings of CIPT, combined with its simplicity, would encourage many Internet entities to engage in CIPT partnerships.

We conclude the paper with several open problems that are the focus of our current investigation:

Open Problem 1. How do changes in CIPT, both in terms of the coalition structure and volume/temporal effects, affect its dynamic?

Open Problem 2. Can we quantify the factors (size, social, market, geography) that influence the CIPT coalition formation process?

Open Problem 3. Shapley value is an implicit metric: it depends not only on the player's behavior but also on the behavior of the other partners in the CIPT. Can we derive more explicit metrics that would approximate the Shapley value closely, while being explicit and simple to calculate?

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

Table 3: The *cosine*-similarity between the transit (T) and peering (P) time series (both downstream and upstream directions).

APPENDIX

A. RELATION OF TRANSIT TO PEERING TRAFFIC

Here we discuss the relationship of the transit and peering traffic in two academic ISPs that publish their network load information: HEANET and SANET. In Figure 9, we depict the peering and transit traffic for both ISPs on Thursday, 13th Jan. 2011. One can observe that the peering and transit traffic profiles are rather similar. To quantify the similarity of the demand patterns we use the *cosine*-similarity between the corresponding demand time series: $X = (x_1, \ldots, x_T)$ and $Y = (y_1, \ldots, y_T)$:

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

The value of $sim(X,Y)$ is equal to the cosine of the angle between the vectors X and Y in the T-dimensional euclidian space. Thus $sim(X,Y) = 1$ if $X = \alpha Y$ for a scalar α; otherwise $sim(X,Y) < 1$. Table 3 reports the values of *cosine*-similarity for the upstream and downstream time series for the both ISPs.

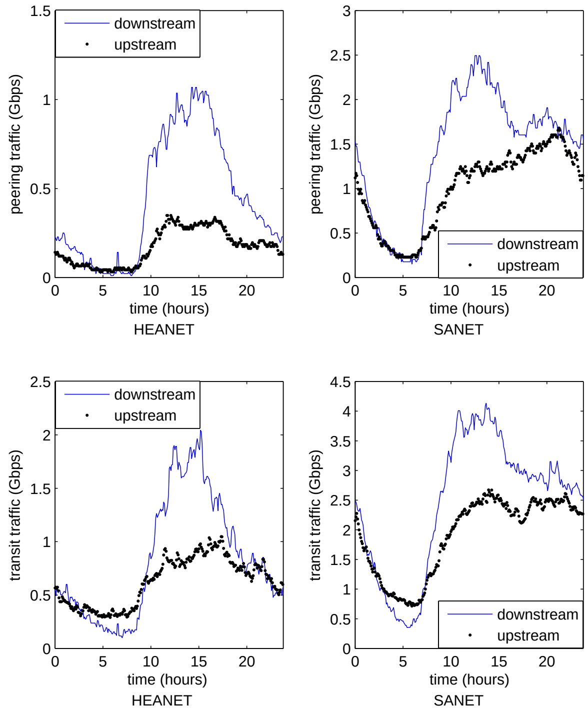

Figure 9: The transit and peering traffic in two national ISPs: HEANET and SANET.

COMMENT 2. *We do not report the statistics from the other two ISPs mentioned in Section 5.1.2, CESNET [12] and GRNET [31], because their visual* rrdtool *images were very nonstandard and our OCR tool could not extract numeric data from them. However, simple visual check can confirm that the transit-peering relationships in these two networks are very similar to those observed in HEANET and SANET.*

B. REFERENCES

[1] A. Akella, B. Maggs, S. Seshan, A. Shaikh, R. Sitaraman, "A Measurement-based Analysis of Multihoming", Proc. of ACM SIGCOMM 2003.
[2] A. Akella, J. Pang, B. Maggs, S. Seshan, A. Shaikh, "A Comparison of Overlay Routing and Multihoming Route Control", Proc. of ACM SIGCOMM 2004.
[3] C. Amega-Selorm, M. Mureithi, D. Pater, R. Southwood, "Impact of IXPs - A review of the experiences of Ghana, Kenya and South Africa", Open Society Institute, 2009.
[4] A. Archer, J. Feigenbaum, A. Krishnamurthy, R. Sami, S. Shenker, "Approximation and Collusion in Multicast Cost Sharing", Games and Economic Behavior, 47(1), 2004.
[5] B. Augustin, B. Krishnamurthy, W. Willinger, "IXPs: Mapped?", Proc. of ACM SIGCOMM 2009.
[6] T. Ballardie, P. Francis, J. Crowcroft, "Core Based Trees (CBT)", Proc. of ACM SIGCOMM 1993.
[7] G. Biczok, T. A. Trinh, "A User-Influenced Pricing Mechanism for Internet Access", In Proc. of ICQT 2009.
[8] V. Bonneau, "Internet Traffic and Economics", DigiWorld Summit 2010.
[9] J. W. Byers, J. Considine, M. Mitzenmacher, S. Rost, "Informed Content Delivery Across Adaptive Overlay Networks", IEEE/ACM Transactions on Networking, 12(5), 2004.
[10] J. W. Byers, M. Luby, M. Mitzenmacher, "A Digital Fountain Approach to Asynchronous Reliable Multicast", IEEE J-SAC, 20(8), 2002.
[11] The Multi Router Traffic Grapher. `http://oss.oetiker.ch/mrtg/`.
[12] CESNET, Traffic Stats. `http://www.ces.net/netreport/`.
[13] K. Cho, K. Fukuda, H. Esaki, A. Kato, "The Impact and Implications of the Growth in Residential User-to-User Traffic", Proc. of ACM SIGCOMM 2006.
[14] D. R. Choffnes, F. E. Bustamante, "Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in Peer-to-peer Systems", Proc. of ACM SIGCOMM 2008.
[15] Cisco Systems, "Cisco Visual Networking Index: Forecast and Methodology", Cisco White Paper, 2008.
[16] D. Clark, "Combining Sender and Receiver Payments in the Internet", Proc. of TRPC 1996.
[17] D. Clark, J. Wroclawski, K. R. Sollins, R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", Proc. of ACM SIGCOMM 2002.
[18] J. Crowcroft, "Net Neutrality: The Technical Side of the Debate ~ A White Paper", ACM SIGCOMM CCR, 37(1), 2007.
[19] S. Deering, "Multicast Routing in a Datagram Internetwork", Ph.D. Thesis, Stanford University, 1991.
[20] A. Dhamdhere, C. Dovrolis, "Can ISPs be Profitable Without Violating Network Neutrality?", Proc. of ACM NetEcon 2008.
[21] A. Dhamdhere, C. Dovrolis, "The Internet is Flat: Modeling the Transition from a Hierarchy to a Peering Mesh", Proc. of ACM CoNEXT 2010.
[22] A. Dhamdhere, C. Dovrolis, P. Francois, "A Value-based Framework for Internet Peering Agreements", Proc. of ITC 2010.
[23] X. Dimitropoulos, P. Hurley, A. Kind, M. Stoecklin, "On the 95-Percentile Billing Method", Proc. of PAM 2009.
[24] M. Dischinger, M. Marcon, S. Guha, K. P Gummadi, R. Mahajan, S. Saroiu, "Glasnost: Enabling End Users to Detect Traffic Differentiation", Proc. of NSDI 2010.
[25] F. Douglis, A. Feldmann, B. Krishnamurthy, J. Mogul, "Rate of Change and other Metrics: a Live Study of the World Wide Web", Proc. of USITS 1997.
[26] European Internet Exchange Association, `www.euro-ix.net`.
[27] J. Feigenbaum, C. Papadimitriou, S. Shenker, "Sharing the Cost of Multicast Transmissions", Computer and System Sciences: Special Issue on Internet Algorithms, 63 (1), 2001.
[28] B. Fletcher, Renesys Official Blog, `http://www.renesys.com/blog/2010/10/internet-transit-sales-2005-10.shtml`, 2010.
[29] Global Internet Geography, Telegeography Technical Report, `http://www.telegeography.com/product-info/gig/download/telegeography-global-internet.pdf`, 2010.
[30] D. K. Goldenberg, L. Qiu, H. Xie, Y. R. Yang, Y. Zhang, "Optimizing Cost and Performance for Multihoming", Proc. of ACM SIGCOMM 2004.
[31] GRNET, Traffic Stats. `http://mon.grnet.gr/rg/`.
[32] HEANET, Traffic Stats. `http://www.hea.net/`.
[33] M. Hrybyk, "The Transit Exchange - A New Model for Open, Competitive Network Services", Proc. of PTC 2007.
[34] P. Hui, R. Mortier, K. Xu, J. Crowcroft, V. O. K. Li, "Sharing Airtime with Shair Avoids Wasting Time and Money", Proc. of Hotmobile 2009.
[35] Information Sciences Institute, University of Southern California, "Internet Protocol DARPA Internet Program Protocol Specification", RFC 791, 1981.

[36] E. Kreifeldt, IT Business Edge Interview, http://www.itbusinessedge.com/cm/community/features/interviews/blog/ip-wholesale-volume-up-costs-down/?cs=44377&page=2, 2010.

[37] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, F. Jahanian, "Internet Inter-domain Traffic", Proc. of ACM SIGCOMM 2010.

[38] N. Laoutaris, G. Smaragdakis, P. Rodriguez, R. Sundaram, "Delay Tolerant Bulk Data Transfers on the Internet", Proc. of ACM SIGMETRICS 2009.

[39] C. Li, K. Sycara, A. Scheller-Wolf, "Combinatorial Coalition Formation for Multi-item Group-buying with Heterogeneous Customers", Decision Support Systems, 49(1), 2010.

[40] D. Liben-Nowell, A. Sharp, T. Wexler, K. Woods, "Computing Shapley Value in Cooperative Supermodular Games", Preprint, 2010.

[41] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, D. Rubenstein, "On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers", IEEE/ACM Transactions on Networking, accepted for publication, 2011.

[42] R. Mahajan, D. Wetherall, T. Anderson, "Negotiation-Based Routing Between Neighboring ISPs", Proc. of NSDI 2005.

[43] M. Marcon, M. Dischinger, K. P. Gummadi, A. Vahdat, "The Local and Global Effects of Traffic Shaping in the Internet", Proc. of COMSNETS 2011.

[44] Minnesota Internet Traffic Studies (MINTS), http://www.dtc.umn.edu/mints, 2009.

[45] V. Misra, S. Ioannidis, A. Chaintreau, L. Massoulie, "Incentivizing Peer-Assisted Services: A Fluid Shapley Value Approach", Proc. of ACM SIGMETRICS 2010.

[46] H. Moulin, S. Shenker, "Strategyproof Sharing of Submodular Costs: Budget Balance versus Efficiency", Economic Theory, 18 (3), 2001.

[47] V. N. Padmanabhan, J. C. Mogul, "Improving HTTP Latency", Computer Networks and ISDN Systems, 28(1-2), 1995.

[48] M. Piatek, A. Krishnamurthy, A. Venkataramani, R. Yang, D. Zhang, A. Jaffe, "Contracts: Practical Contribution Incentives for P2P Live Streaming", Proc. of NSDI 2010.

[49] L. Qiu, V. N. Padmanabhan, G. M. Voelker, "On the Placement of Web Server Replicas", Proc. of IEEE INFOCOM 2001.

[50] Y. Rekhter, T. Li, "A Border Gateway Protocol (BGP-4)", RFC 1771, 1995.

[51] M. Roussopoulos, M. Baker, D. S. H. Rosenthal, T. J. Giuli, P. Maniatis, J. Mogul, "2 P2P or Not 2 P2P?", Proc. of IPTPS 2004.

[52] SANET, Traffic Stats. http://samon.cvt.stuba.sk/.

[53] L. S. Shapley, "A Value for n-Person Games", Annals of Mathematics Study, 28, 1953.

[54] F. B. Shepherd, G. Wilfong, "Multilateral Transport Games", Proc. of INOC 2005.

[55] G. Shrimali, A. Akella, A. Mutapcic, "Cooperative Interdomain Traffic Engineering Using Nash Bargaining and Decomposition", IEEE/ACM Transactions on Networking, 18(9), 2007.

[56] G. Smaragdakis, V. Lekakis, N. Laoutaris, A. Bestavros, J. W. Byers, M. Roussopoulos, "EGOIST: Overlay Routing using Selfish Neighbor Selection", Proc. of ACM CoNEXT 2008.

[57] J. Sommers, P. Barford, N. Duffield, A. Ron, "Multi-objective Monitoring for SLA Compliance", IEEE/ACM Transactions on Networking, 18(2), 2010.

[58] R. Stanojevic, N. Laoutaris, P. Rodriguez, "On Economic Heavy Hitters: Shapley Value Analysis of the 95th-Percentile Pricing", Proc. of ACM IMC 2010.

[59] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, H. Balakrishnan, "A Scalable Peer-to-peer Lookup Protocol for Internet Applications", IEEE/ACM Transactions on Networking, 11(1), 2003.

[60] Voxel dot net. IP-services pricelist. https://www.voxel.net/ip-services

[61] E. Winter, "The Shapley Value", The Handbook of Game Theory, North-Holland, 2002.

[62] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, A. Silberschatz, "P4P: Provider Portal for Applications", Proc. of ACM SIGCOMM 2008.

[63] Y. Zhu, C. Dovrolis, M. Ammar, "Combing Multihoming with Overlay Routing (or How to Be a Better ISP Without Owning a Network)", Proc. of IEEE INFOCOM 2007.