

# CIPT: Using Tuangou to Reduce IP Transit Costs\*

Rade Stanojevic<sup>†</sup>   Ignacio Castro   Sergey Gorinsky  
Institute IMDEA Networks, Madrid, Spain

## 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 of 264 ISPs 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.

## Categories and Subject Descriptors

C.2.3 [Computer-communications networks]: Network operations; network management

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

<sup>†</sup>R.S. is also an adjunct professor at Universidad Carlos III de Madrid.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

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

Copyright 2011 ACM 978-1-4503-1041-3/11/0012 ...\$10.00.

## General Terms

Economics

## Keywords

Network economics; burstable billing; group buying; cost sharing; Shapley value

## 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) 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 [21], the interdomain traffic grows [27], and the overall IP transit costs remain high or even increase according to industry analysts [6, 20, 26].

The problem of reducing the IP transit costs has attracted notable solutions of IXPs (Internet eXchange Points) [4, 14], IP multicast [13], CDNs (Content Distribution Networks) [38], P2P (Peer-to-Peer) localization [10], and traffic smoothing [28, 34]. 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 [29]. Similarly to the tuangou elsewhere, CIPT succeeds primarily due

to subadditivity of prices [21, 44, 45]. However, the benefits of CIPT depend also on burstable billing [16], different methods to account for bidirectional traffic, and other complex factors.

Relying on real inter-domain traffic and transit pricing, we empirically estimate the potential gains from CIPT and tradeoffs between several factors that affect these gains. We also propose to use Shapley value [46] 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 (*mrtg* [7]) traffic statistics available 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 (national) 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 benefits of CIPT.

The two critical properties that enable CIPT, namely the subadditive pricing and 95-th percentile billing, are ubiquitous in the IP transit market and have existed from early days of the commercial Internet. The question of how long the IP pricing retains these properties remains open. In the entertainment industry (cinema, sports, etc.) that have offered group discounts for decades, group-discounted tickets still constitute a minor fraction of the revenues. Although IP transit is different from the entertainment industry in many ways, we believe in a similar outcome for CIPT: CIPT coalitions will not dominate the IP transit market nor change the transit pricing structure. However, if CIPT coalitions become very common or large, their emergence can lead to a modified pricing structure and other new strategic behaviors in the Internet. An in-depth analysis of this phenomenon is beyond the scope of this paper.

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

**Table 1: IP transit pricing rates of Voxel.net.**

While our results on the CIPT cost reduction validate the potential of CIPT to become a new element of the Internet ecosystem, the practical viability of CIPT also depends on other strategic and organizational issues. These issues are discussed in Section 6 and, in more depth, in the technical report [43].

## 2. BACKGROUND AND MOTIVATION

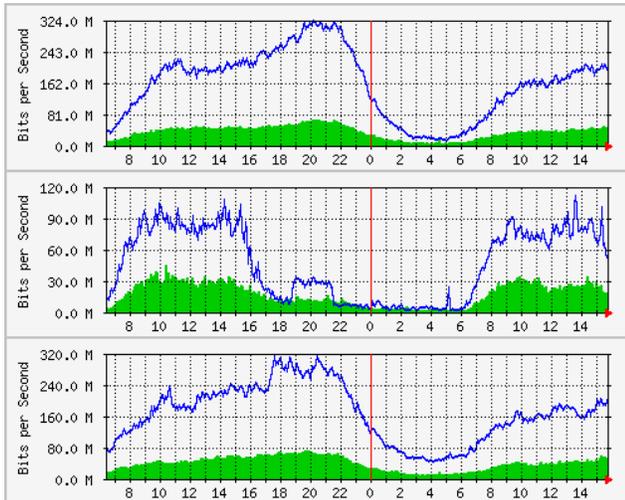
The geographic location affects significantly the cost of IP transit. The IP transit prices per *Mbps* per month range usually from \$5 to \$100 (we use \$, US\$ or USD to refer to U.S. dollars throughout the paper): the wholesale IP transit is typically priced under \$10 per *Mbps* in most European and North American hubs but can exceed \$100 per *Mbps* in Australia, Latin America and other remote regions of the Internet [2, 21].

Regardless of the geographic location, IP transit is subject to economies of scale and is priced subadditively: the prices per *Mbps* per month<sup>1</sup> are smaller for larger quantities of IP transit [21, 45, 44]. Table 1 presents the current (as of June 2011) transit pricing rates of Voxel dot net [45]. 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 pay \$5000.

Burstable billing is another important aspect of IP transit pricing [16, 34]. To calculate the IP transit cost, the most commonly used method is to calculate the *peak* usage (typically through the 95th-percentile rule [16, 34]) and then the price function  $f$  is applied to 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 CIPT because two ISPs with their traffic peaks in opposite directions can mutually benefit from the less utilized direction of each other.

Finally, it is important noting that the prices (per

<sup>1</sup>Throughout the paper we use the cost unit \$ per *Mbps*, omitting the charging period: one month.



**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.**

$Mbps$ ) of transit are in constant decline over the previous decade, with an average decay of around 25 – 30% per year [21]. Even though market forces inevitably drive this trend, the total amount of interdomain traffic grows with a rate that outpaces the decay in prices. The recent paper from Arbor Networks [27] reports an annualized inter-domain traffic growth (on a set of 110 geo-diverse ISPs) of 44.5%. CISCO [11] and MINTS [35] report slightly higher annual growth figures, in the range of 50 – 60%. While a fair fraction of this growth is due to increased peering [4], 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” [26]. Similar observations have been made by several other business analysts [6, 20].

To illustrate the potential of CIPT, we consider a simple scenario of three partners<sup>2</sup>  $P_1$ ,  $P_2$ , and  $P_3$  interested in purchasing IP transit from the same provider. We assume the transit pricing rates as in Table 1, 95th-percentile burststable 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

<sup>2</sup>We interchangeably use terms partner, subject and player to refer to an ISP, hosting provider, website or any other entity interested in purchasing IP transit.

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 (per month).

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. Burststable 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 decrease even if the pricing function were additive;
2. Subadditive pricing – the upgrade from the 100- $Mbps$  CDR level to the 1000- $Mbps$  one yields a lower price per  $Mbps$  and thereby reduces the aggregate transit cost even further.

### 3. COOPERATIVE IP TRANSIT

In the previous sections we hinted at the main idea of the CIPT. This section provides more details and discusses several aspects of the strategy.

We use the term *Cooperative IP Transit* (CIPT) to refer to any cooperative scheme 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 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  $i \in \{1, 2, \dots, N\}$ , and time  $t$  is measured in fixed-size time intervals with a typical interval duration of 5 minutes. The cost that 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. 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 (see later sections for empirical validation of the fact):

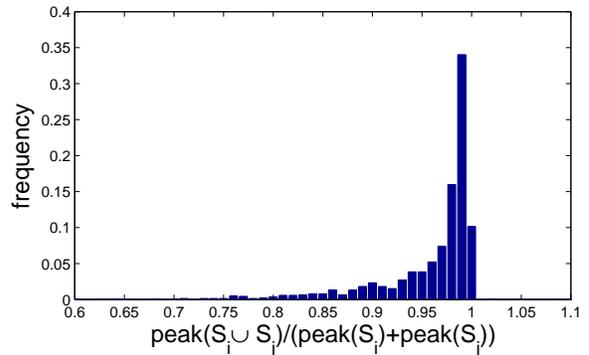
$$\rho = \frac{C}{\sum_{i=1}^N C_i} < 1.$$

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 results in 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.

### 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 par-



**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 Slovak IXP.**

titive<sup>3</sup> set of  $\mathcal{N}$  to a cost value:  $c : 2^{\mathcal{N}} \rightarrow R$ . 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 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 in which the sum of the 95th-percentiles is smaller than the 95th-percentile of the union of the traffic of the two subjects<sup>4</sup>. 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 Slovak Internet Exchange (SIX). The SIX and several other IXPs publish statistics for the traffic exchanged by each of their members (mostly regional ISPs). This information represents a valuable proxy for estimating the traffic patterns (volume, peak-hour, peak-to-valley ratio, up/downstream

<sup>3</sup>For set  $\mathcal{N}$ , the partitive set of  $\mathcal{N}$  is the set of all subsets of  $\mathcal{N}$  and is usually denoted as  $2^{\mathcal{N}}$ .

<sup>4</sup>For example, two subjects consuming 100  $Mbps$  4% of the time each, one in the morning and the other overnight, 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  $Mbps$ .

traffic ratio, etc.) for the involved ISPs; see Section 5.1.2 and Appendix for more details.

*Observation 1.* *The traffic patterns of potential CIPT members 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 cost of cooperation. As we saw in Section 3.1 the CIPT can be abstracted as a cooperative game which allows us 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, kernel, nucleolus, and Shapley value [46]. While other solution concepts have attractive features, in the context of CIPT we find particularly appealing to use Shapley value since it has several distinct important properties, i.e., Shapley value: (1) *exists* for any cooperative game and is *uniquely* determined, (2) satisfies basic *fairness* postulates [40, 46], 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 coalition.

For a given cooperative game  $(\mathcal{N}, c(\cdot))$ , the Shapley value is a (unique) vector  $(\phi_1(c), \dots, \phi_N(c))$  defined below, for sharing the cost  $c(\mathcal{N})$  generated by 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 [40, 46] 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:

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

where the sum is taken across all permutations (or arrival orders)  $\pi$  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  $\pi$ . 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 Eq. (2), it is not practically feasible to employ Eq. (2) for  $N > 30$ . A number of methods have been suggested for accurate estimation of Shapley value, and in this paper we use a simple Monte Carlo method, analyzed in [30], 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  $\Pi_k$  of  $K$  randomly sampled arrival orders:

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

The parameter  $K$  determines the error between the real Shapley value and its estimate: the greater  $K$  is the lower the error becomes. Thus, 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 expected error of under 1% across all the CIPT players. Therefore, 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. We also calculate the breakdown of the impact of the two complementary properties of IP transit pricing, and conclude that while most of the savings can be expected from the subadditivity of prices, a non-negligible share of the savings comes from the 95-th percentile billing<sup>5</sup>. In Section 5.4 we augment this anal-

<sup>5</sup>For geo-diverse coalitions the share of savings coming from

IXP	acronym	# 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
Slovak IX	SIX	52	42	23	15.4%/44.9%	0.27
Israeli IX	IIX	17	2.1	1.38	14.3%/40.6%	0
Finnish 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.

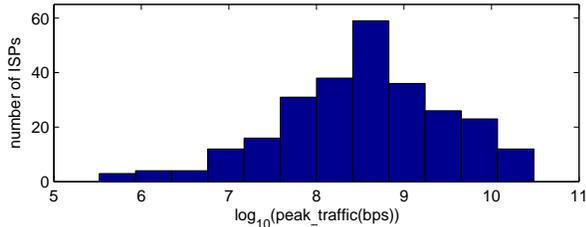


Figure 3: The distribution of the peak traffic rates across all 264 ISPs: median 560 Mbps; mean 2.9 Gbps.

ysis 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 5.3 analyzes the per-partner savings and shows somewhat expectable trends that the larger the partner is, the larger are its absolute savings, but the smaller are its relative savings. In the technical report [43], we present additional empirical results, including the analysis of the effects of collaboration between geo-diverse partners 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 (IXPs) report upstream and downstream demand time series for the traffic exchanged by every member of the IXPs. Those that do it are listed in Table 2. This traffic statistics data is typically given in the form of `mrtg` images [7], 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. Section 5.1.2 considers a small set of ISPs that make their detailed traffic information public, and presents an empiric argument showing that the IXP traffic of an ISP is a good proxy for estimating its transit, 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 95th-percentile billing is likely to be more pronounced.

the medium-sized and large IXPs [18]. A majority of these IXPs publish their aggregate traffic statistics, summed across all the members, but some, including those listed in Table 2, also make public the detailed traffic statistics of their members. We then crawled the websites of these 6 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` [7]. To convert the information into a numeric form, we built a piece of software that takes as input an `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. Raw visual data, the numeric data and the code for transforming `mrtg/rrdtool` images into the numeric format can be found at [12].

### 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 a detailed picture of both their network infrastructure and utilization of their networks. Those that we identified are HEANET (Ireland) [23], SANET (Slovak Republic) [39], CESNET (Czech Republic) [8], GRNET (Greece) [22]. 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. In those 4 ISPs, peering corresponds to 35-40% of the total traffic, with the remaining 65-60% 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 the Appendix 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)




Figure 4: The absolute/relative savings as a function of the transit to IXP traffic ratio.

with their peering traffic (information that is publicly available) multiplied by a factor  $\gamma$  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  $\gamma \in [0.5, 4]$ . In Sections 5.3 and 5.4 (which analyze the cost-sharing and coalition size), we fix  $\gamma = 1.5$ , that corresponds to transit vs. peering traffic ratio of 60 : 40 as suggested by our analysis in the 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 present the results for both the sum and max models. In Sections 5.3 and 5.4 we focus on the sum pricing model (the more conservative one in terms of cost reduction) for the analysis of cost-sharing and coalition size.

## 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 the IXPs listed in Table 2. Following the discussion in Section 5.1.2, we approximate the IP transit traffic patterns with the traffic exchanged

at these IXPs multiplied by a constant  $\gamma \in [0.5, 4]$ ; this constant represents the ratio between the transit and IXP traffic volume. As we discussed above, this approximation 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 [9] 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 non-local IP markets [2]). The empirical evidence of the few European ISPs discussed in the Appendix suggests that, in medium-sized European countries with one dominant IXP,  $\gamma$  is in the range between 1.5 and 2.

We stress again that the purpose of this evaluation is to shed some light on 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 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. These 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 Table 2 that for the IXPs with a low *skewness* of under 0.3 (SIX, IIX, and Inter-LAN), the expected relative savings are considerably higher than those of the others with high *skewness*.


Remember that the savings of CIPT come from two properties of the IP transit model: price subadditivity and 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 subadditivity of the prices, i.e., if the price per *Mbps* would be constant independent of the usage level. Such savings would come exclusively from the reduction in the 95th-percentile. Table 2 reports these results in the column called *95th-pct effect*. From this table we can conclude that both properties (price subadditivity and the 95th-percentile billing) influence the total savings by a non-trivial amount. However, the exact breakdown of the impact of these two properties on the total savings depends on the volume of the traffic and 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 a low per-*Mbps* price). The relative savings are, however, bounded from below by the *95th-pct effect* quantity reported in Table 2 for both sum and max 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 the Appendix). Using the corresponding value  $\gamma \approx 1.5$ , 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 Per-partner 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  $\gamma$  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 5 against the original IP transit annual costs. Figure 6 shows the absolute annual savings (the difference between the original IP



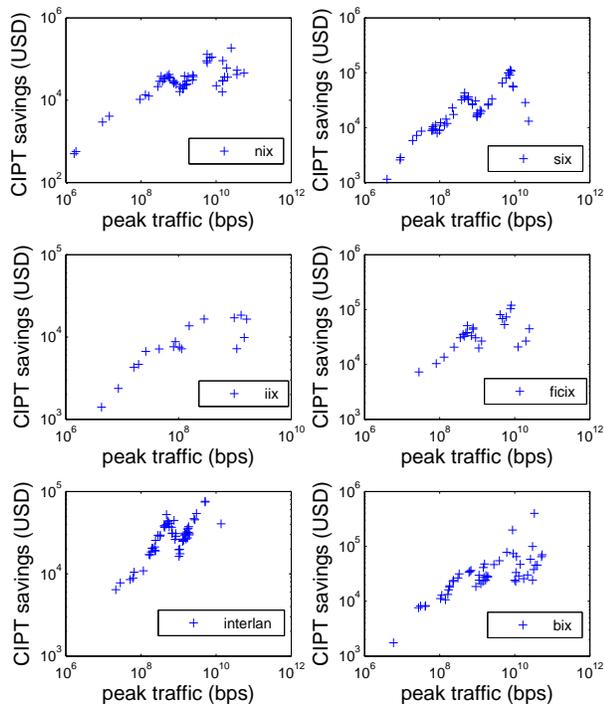
**Figure 5: The original annual costs versus CIPT costs (Shapley value) across all the ISPs from the 6 IXPs.**

transit costs and CIPT costs) for all ISPs in these 6 IXPs.

We can observe two trends in Figures 5 and 6. 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* prices, 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, 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 and to the fact that the peak-hour of the coalition is predominantly determined by large ISPs. This means that large ISPs that join already large-enough coalitions (those that reach a close-to-minimum price per *Mbps*) do not bring great benefits to the coalition and consequently implying low (yet positive) relative gains for these ISPs.

### 5.4 Coalition size

In the 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. Namely, the managing costs of the coalition grow with the coali-

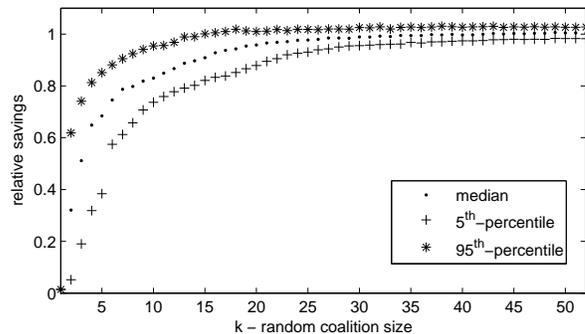


**Figure 6: The absolute annual savings for all the ISPs from the 6 IXPs.**

tion size. Additionally, some ISPs may be unwilling to participate in a common business arrangement or prevented from the participation by their current contract with the transit provider. See [43] for in depth discussion on the CIPT organizational embodiment.

In this section, we show that small coalitions can offer savings comparable to those of the large coalitions. For that reason, we take the Slovak IXP (SIX) with  $N = 52$  members, and for each  $k \in \{1, 2, \dots, N\}$  we analyze the per-partner savings from participating in the coalition of  $k$  random members of SIX. We use the `sum` pricing model and  $\gamma = 1.5$ . The results for other IXPs, `max` pricing model and other choices of  $\gamma$  are very similar, hence we omit them for brevity.

In Figure 7 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 7 we can observe the law of diminishing returns: relatively small coalitions provide savings very close to the savings of the large coalitions, and, adding more members to the coalition decreases the incremental savings. In particular, even with as few as  $k = 3$



**Figure 7: Relative (as fraction of the savings obtained in the grand coalition) per-partner savings for smaller coalitions.**

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 with 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, an ISP can obtain even higher savings, as the 95th-percentile of the savings in Figure 7 suggests. However, such optimization is out of scope for the present paper.

## 6. CIPT BEYOND THE COST SHARING

Section 5.3 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 and strategic issues.

**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 technical and management overhead costs of operating the CIPT and expected to be insignificant in comparison to the transit cost reductions provided by the CIPT. In a future study, we plan to quantify the technical and economic overhead. 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. Finally, a single ISP may choose to participate in multiple CIPTs in order to increase the provider diversity.

**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 [24]. 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) [45] 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) 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.

**Transit providers and strategic issues:** The costs saved by CIPT coalitions are not necessarily the revenues lost by transit providers due to at least two factors: aggregation of small traffic rates and competition among transit providers. Typically, a large ISP is not interested in selling IP transit directly to small ISPs. Instead, a middle-sized ISP resells the IP transit to the small customers. By forming a CIPT, the small ISPs can make their aggregate traffic rate large enough for the large provider and bypass the middleman. On the other hand, the large ISP finds in the CIPT a new customer. 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. While a transit provider might be inclined to offer a CIPT a contract with higher prices than in its standard individual contract, this strategy is dangerous for the provider because IP transit is a competitive market: the CIPT can take its transit business to another ISP, e.g., an ISP that offers only one type of transit contracts. In general, we do not expect CIPT coalitions to dominate the IP transit market or change the transit pricing structure. However, if CIPT coalitions become very common or large, this can lead to the common use of CIPT-type contracts with higher prices. Our technical report [43] explores these and other strategic issues in more detail.

## 7. RELATED WORK

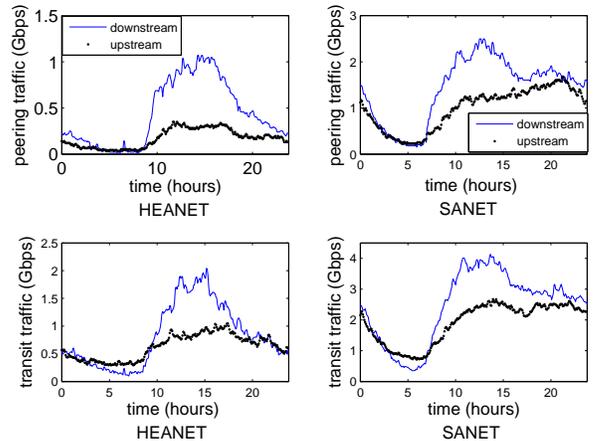
In presenting and evaluating CIPT, we already mentioned the essential background information. This section takes a broader look at related work.

The existing approaches for reducing the transit costs include ISP peering, IP multicast, CDNs, P2P localization, and traffic smoothing. Peering [4, 14] 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 [13] duplicates packets in IP routers and thereby reduces transit traffic. While IP multicast requires router support from transit providers, CDNs [38] 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 [10] 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 [16]. An ISP can do such traffic smoothing with rate limiting [34] or in-network storage for delay-tolerant traffic [28]. Unlike the above approaches that modify the transit traffic, CIPT reduces the transit costs without altering it. Most recently, from the transit provider perspective, [44] argues that having a small number of different pricing tiers (different price levels for varying bandwidth commit levels) is enough for maximizing the IP transit provider profits.

We view CIPT as a coalition and use the Shapley value [40] for sharing CIPT costs. Shair [25] 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 [3, 19] and interdomain routing [33, 41]. The game-theoretic analyses of the Shapley-value mechanism [3, 19, 37] 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 [31, 32], incentives in P2P systems [36], and charging individual users by access ISPs [42] 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 [30].

## 8. CONCLUSIONS

In spite of the steady decline of IP transit prices, 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.



**Figure 8: The transit and peering traffic in two national ISPs: HEANET and SANET.**

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. Instead, CIPT utilizes *tuangou*, or team-buying, for IP transit. The savings in CIPT come from two distinct yet ubiquitous properties of the IP transit pricing model: price subadditivity 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 affect its dynamic?

*Open Problem 2.* Can we quantify the factors that influence the CIPT coalition formation process?

*Open Problem 3.* Can we derive more suitable metrics that would approximate the Shapley value closely, while being explicit and simple to calculate?

*Open Problem 4.* What would be the effect of CIPT on the Internet AS-level topology?

## Acknowledgments

This work was supported in part by Department-of-Education grant S2009/TIC-1468 from the Regional Government of Madrid, Ramon-y-Cajal grant RYC-2009-04660 from the Spanish Ministry of Science and Innovation, FP7-PEOPLE grant 229599 and FP7-ICT grant 288021 from the European Commission.

## 9. REFERENCES

- [1] P. Antoniadis, S. Fdida, T. Friedman, V. Misra, “Federation of Virtualized Infrastructures: Sharing the Value of Diversity”. Proc. of ACM CoNEXT 2010.

[2] C. Amega-Selorm, et al., “Impact of IXPs - A Review of the Experiences of Ghana, Kenya and South Africa”, Open Society Institute, 2009.

[3] A. Archer et al., “Approximation and Collusion in Multicast Cost Sharing”, Games and Economic Behavior, 47(1), 2004.

[4] B. Augustin, B. Krishnamurthy, W. Willinger, “IXPs: Mapped?”, Proc. of ACM SIGCOMM 2009.

[5] G. Biczok, T. A. Trinh, “A User-Influenced Pricing Mechanism for Internet Access”, In Proc. of ICQT 2009.

[6] V. Bonneau, “Internet Traffic and Economics”, DigiWorld Summit 2010.

[7] The Multi Router Traffic Grapher. <http://oss.oetiker.ch/mrtg/>.

[8] CESNET, Traffic Stats. <http://www.ces.net/netreport/>.

[9] 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.

[10] 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.

[11] Cisco Systems, “Cisco Visual Networking Index: Forecast and Methodology”, Cisco White Paper, 2008.

[12] CIPT dataset: [http://fourier.networks.imdea.org/people/~rade\\_stanojevic/CIPTdata/](http://fourier.networks.imdea.org/people/~rade_stanojevic/CIPTdata/)

[13] S. Deering, “Multicast Routing in a Datagram Internetwork”, Ph.D. Thesis, Stanford University, 1991.

[14] A. Dhamdhere, C. Dovrolis, “The Internet is Flat: Modeling the Transition from a Hierarchy to a Peering Mesh”, Proc. of ACM CoNEXT 2010.

[15] A. Dhamdhere, C. Dovrolis, P. Francois, “A Value-based Framework for Internet Peering Agreements”, Proc. of ITC 2010.

[16] X. Dimitropoulos, P. Hurley, A. Kind, M. Stoeklin, “On the 95-Percentile Billing Method”, Proc. of PAM 2009.

[17] M. Dischinger et al., “Glasnost: Enabling End Users to Detect Traffic Differentiation”, Proc. of NSDI 2010.

[18] European Internet Exchange Association, [www.euro-ix.net](http://www.euro-ix.net).

[19] J. Feigenbaum, C. Papadimitriou, S. Shenker, “Sharing the Cost of Multicast Transmissions”, Computer and System Sciences: Special Issue on Internet Algorithms, 63 (1), 2001.

[20] B. Fletcher, Renesys Official Blog, <http://www.renesys.com/blog/2010/10/internet-transit-sales-2005-10.shtml>, 2010.

[21] Global Internet Geography, Telegeography Technical Report, <http://www.telegeography.com/product-info/gig/download/telegeography-global-internet.pdf>, 2010.

[22] GRNET, Traffic Stats. <http://mon.grnet.gr/rg/>.

[23] HEANET, Traffic Stats. <http://www.heanet.net/>.

[24] M. Hrybyk, “The Transit Exchange - A New Model for Open, Competitive Network Services”, Proc. of PTC 2007.

[25] 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.

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

[27] C. Labovitz et al., “Internet Inter-domain Traffic”, Proc. of ACM SIGCOMM 2010.

[28] N. Laoutaris et al., “Delay Tolerant Bulk Data Transfers on the Internet”, Proc. of ACM SIGMETRICS 2009.

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

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

[31] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, D. Rubenstein, “Internet Economics: The use of Shapley value for ISP settlement”, IEEE/ACM Transactions on Networking, vol. 18(3), 2010.

[32] 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, vol. 19(3), 2011.

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

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

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

[36] V. Misra, S. Ioannidis, A. Chaintreau, L. Massoulié, “Incentivizing Peer-Assisted Services: A Fluid Shapley Value

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).**

Approach”, Proc. of ACM SIGMETRICS 2010.

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

[38] L. Qiu et al., “On the Placement of Web Server Replicas”, Proc. of IEEE INFOCOM 2001.

[39] SANET, Traffic Stats. <http://samon.cvt.stuba.sk/>.

[40] L. S. Shapley, “A Value for  $n$ -Person Games”, Annals of Mathematical Studies, 28, 1953.

[41] G. Shrimali, A. Akella, A. Mutapcic, “Cooperative Interdomain Traffic Engineering Using Nash Bargaining and Decomposition”, IEEE/ACM Trans. on Netw., 18(9), 2007.

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

[43] R. Stanojevic, I. Castro, S. Gorinsky, “CIPT: Using Tuangou to Reduce IP Transit Costs”, Technical Report TR-IMDEA Networks-2011-4, Institute IMDEA Networks, 2011.

[44] V. Valancius, et al., “How Many Tiers? Pricing in the Internet Transit Market”, Proc of ACM SIGCOMM, 2011

[45] Voxel dot net, pricelist. <https://www.voxel.net/ip-services>

[46] E. Winter, “The Shapley Value”, The Handbook of Game Theory, North-Holland, 2002.

## APPENDIX: RELATIONSHIP 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 8, 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:  $X = (x_1, \dots, x_T)$  and  $Y = (y_1, \dots, y_T)$ :

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

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  $\alpha$ ; otherwise  $sim(X, Y) < 1$ . Table 3 reports the values of cosine-similarity for the upstream and downstream time series for the both ISPs.

**Remark.** We do not report the statistics from the other two ISPs mentioned in Section 5.1.2, CESNET [8] and GRNET [22], 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.