Can Bilateral ISP Peering Lead to Network-wide Cooperative Settlement

Yang Cheung, Dah-ming Chiu, Jianwei Huang
Department of Information Engineering, The Chinese University of Hong Kong
Email: {ycheung6, dmchiu, jwhuang}@ie.cuhk.edu.hk

***Abstract*— The Internet includes thousands of Internet service providers (ISPs) which are interconnected to provide connectivity and service for end-users. Traditionally, the settlement between the ISPs are determined based on bilateral agreements that result from pair-wise negotiations. Although this settlement mechanism is intuitive and easy to implement, it does not encourage network-wide cooperation, as the bilateral charges typically do not lead to a fair division of revenue among all ISPs that are involved in carrying the same flows of traffic. This problem is getting more severe with various emerging new Internet business models.**

In this paper, we try to determine the existence and realizability of bilateral prices that can achieve fair revenue division among ISPs. In particular, we use Shapley value as the basis for deriving fair prices. Under a quite general topology and traffic model, we find that there exists prices that make the revenue division under bilateral settlement equal to that calculated under Shapley value. The corresponding "fair price" exhibits several nice and desirable characteristics. Moreover, it could be realized approximately.

I. Introduction

The Internet is the result of interconnecting many networks, each operated by a separate ISP. The Internet service, as enjoyed by the customer of all ISP networks, is however a service provided collectively by all ISPs. A customer of the Internet service subscribes to an access ISP. For this service, the access ISP collects a charge from each subscriber. Since the service is accomplished by multiple ISPs, a fundamental issue is how ISPs should divide up this charge.

Based on the design of the Internet, the unit of network service is a packet transported. This is a rather miniscule unit for measuring service, not to mention the task of dividing the charge among multiple parties contributing to the service. By convention (established historically), ISPs charge users monthly on a flat rate basis (like an all-you-can-eat buffet), and settle account monthly on a bilateral basis between ISPs who are connected with each other. There are two most common types of bilateral peering relationships between ISPs. In the first type of peering relationship, one ISP is considered as a transit provider for the other ISP, and the transit provider charges its customer ISP for amount of traffic transit through the provider network. The second type of peering relationship is a totally collaborative one. Two ISPs exchange traffic and deem to benefit from it mutually and forego any charges to each other. Although this all seems a rather sloppy business practice, it keeps the effort in book-keeping to a minimum. This minimalist approach is also totally decentralized, without the need for all kinds of coordination between ISPs.

In recent years, new business models (using Internet services) emerge; for example, Internet content providers (ICPs) are able to generate significantly higher levels of revenues [1]. This prompted re-examination of how the benefits of network services should be divided between various players. One study [2] specifically re-examined the issue of how ISPs should share the total revenues of the Internet service and proposed Shapley values as a potential solution. Shapley value refers to an axiomatically derived formula for fairly dividing a prize among a set of contributors in a general economic setting. By forming an coalition, it is argued in [2] that the ISPs will find more optimal routes and maximize the overall value of the Internet.

In this paper, we ask a different question: Can the current bilateral peering between ISPs, under proper pricing strategies, lead to a cooperative settlement as if the ISPs are all in a coalition? This proposition is not entirely unreasonable, since ISPs do realize the positive network externality in the interconnected network. There are actually two parts to this question: a) does there exist fair pricing (defined as prices in bilateral peering that produce Shapley value as settlements)? and b) how feasible and likely ISPs will choose such fair prices in their peering agreements. For a general set of ISP network topologies and traffic models, we show the answer to (a) is true. We then show some preliminary results to (b).

The paper is organized as follows. We formulate the system model and settlement model in section II and III. In section IV, we show that in a symmetric network setting, we could obtain the fair prices that produce Shapley value and show some desirable properties of the fair prices in Section V. We further consider how fair prices can be calculated in the more general asymmetric topology in Section VI. In Section VII, we show how the fair prices can be locally approximated without knowing the global network information. We conclude this work in section VIII.

II. System Models

A. Topology Model

We use an AS-level hierarchical model to capture the essence of the current tier-based Internet. The hierarchical model has been extensively used in related literatures (eg. [3], [4]). We assume that each ISP contains only one AS.[1] Figure

[1]For an ISP that contains more than one AS, it is enough to think it as one AS in our model.

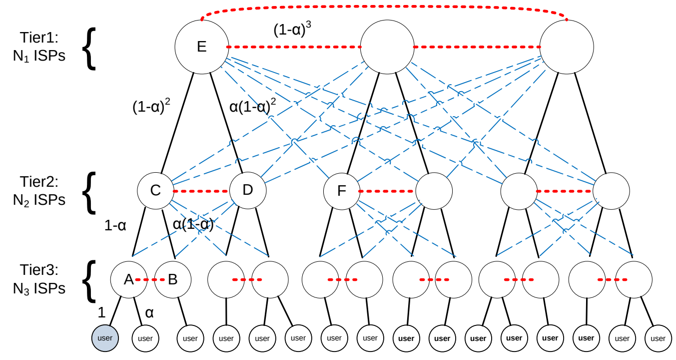

Fig. 1. An illustrative example of the topology and traffic model. Here we have three tiers of ISPs, with three different links: transit links (solid lines), peering links (dotted lines), and multi-homing links (dash lines). For a unit of traffic sent by an end-user of leaf ISP node A (greyed circle), the values beside some of the links represent the probability that the traffic goes to a particular height.

1 is an example of such topology. A set of $\boldsymbol{\mathcal{N}} = \{\mathcal{N}_t, \forall t \in \{1, \dots, T\}\}$ ISPs is located at different tiers, where the set of ISPs in tier t is denoted by $\mathcal{N}_t$. The highest tier (tier 1) is the backbone tier and the lowest tier (tier T, $T = 3$ in Figure 1) connects to end-users. For a particular ISP node $n \in \mathcal{N}_t$ with $t \in \{1, \dots, T-1\}$, it serves as the *provider* for a fixed number of tier $t+1$ ISPs (i.e., *customers*). Different ISPs can have different number of customers. Meanwhile, the same ISP can purchase several different upstream links simultaneously. In this case, we define one of the providers as the *primary provider*. In Figure 1, the primary provider ISP of both node A and node B is node C. For a leaf ISP at tier T, it is the provider for a fixed number of end-users, where each end-user subscribes to one tier T ISP. Furthermore, we assume each ISP has at least two subscribers, so that local routing is always feasible.

We define *region* as the set of nodes which are located at the same tier and share the same primary provider ISP. In Figure 1, node A and node B are in the same region; node D and node F, node A and node C are in different regions.

We classify the links into three types:

1) *Transit link*: a cross-tier link that connects an ISP node to its primary provider ISP at the upper tier. All ISP nodes in the same region establish a transit link to the same provider ISP. In Figure 1, the solid lines represent transit links. Node C and node D are both connected to their common primary provider ISP node E through different transit links.
2) *Peering link*: a link that connects ISPs in the same tier within the same region. Peering links are marked as dotted lines in Figure 1. There is a peering link between node C and node D, but there can not be any peering link between node D and node F since they are not in the same region.
3) *Multi-homing link*: a link that allows an ISP to connect to another upper-tier ISP which is different from but in the same region as its primary provider ISP. In Figure 1, multi-homing links are marked as dash lines. Node A establishes a multi-homing link to node D, but can not establish a multi-homing link to node F since it is not in the same region as its primary provider ISP node C.

Next we define several network topology parameters. First we formulate the peering strategy among ISPs in a macroscopic manner. For all nodes in the same region, we assume that there either does not exist any peering link, or there is a peering link between any pair of nodes such that they form a full mesh topology.[2] We denote $\boldsymbol{\rho} = \{\rho_n, \forall n \in \boldsymbol{\mathcal{N}}\}$ as the set of probabilities that such mesh connection of peering links exists at the region of each node n. Notice that if two nodes n and m are in the same region, then $\rho_n = \rho_m$. Since in practice the topmost tier (tier 1) nodes are typically interconnected into a full mesh through peering links ([3], [5]), we let $\rho_n = 1$ for all nodes $n \in \mathcal{N}_1$.

We denote the multi-homing degree $\boldsymbol{d} = \{d_n, \forall n \in \boldsymbol{\mathcal{N}}\}$ as the number of high-tier ISPs that each node n connects to. When $d_n = 1$, ISP n only connects to its primary provider ISP and thus there is no multi-homing. In Figure 1, $d_C = d_D = 3$ and $d_A = d_B = 2$. We assume that a node is able to use any of its multi-homing links to deliver traffic to all destinations.[3]

In later sections, we will denote a topology $\mathcal{G}$ as $(T, \boldsymbol{\mathcal{N}}, \boldsymbol{\rho}, \boldsymbol{d})$.

B. Traffic Model

We adopt the traffic model used in [4]. In this model, only end-users (who subscribe to the leaf tier T ISPs) are able to initiate and terminate traffic (i.e., act as source and destination). To formulate this user-to-user traffic, we use a parameter α to illustrate how far the traffic is likely to go, instead of specifying the exact destination of the traffic.

For node $n \in \boldsymbol{\mathcal{N}}$, there is a probability α_n for the traffic to go local with intensity I_n. For example in Figure 1, when an end-user of node A has traffic parameter α, his traffic terminates at another end-user of node A with probability α, goes to node C with probability $(1-\alpha)$, and terminates at an end-user of node B with probability $\alpha(1-\alpha)$. The furthest a traffic can go is with probability $(1-\alpha)^T$, which the traffic reaches an end-user that needed to transverse through a peering link at tier 1.

We assume BGP routing is used in the network. In BGP routing, ISPs connected with peering links only exchange routes of its customers to each other. As a result, traffic routes in the hierarchical model have hill-shape appearances (e.g., [6]) - they go uphill until the highest necessary tier, go flat through the peering link if needed, and go downhill to the destination. Because of this, the smallest number of tiers that a traffic reaches essentially determines the distance of that traffic.

[2] The "either none or full mesh" assumption is used to guarantee all traffic that need to go from one ISP node to another ISP node in the same region experience the same hop count: either two-hop by going through the primary provider ISP or one-hop by going through one of the peering links.

[3] This assumption is used to simplify the calculation of Shapley value. In reality, however, there may be some destinations that only some of the multi-homing links can reach.

III. Settlement Model and Definition of Fair Price

Next we describe the two settlement mechanisms to be studied in this paper, namely the bilateral peering settlement model (simplified as *Bilateral*) and the Shapley value cooperative settlement model (simplified as *Shapley*).

A. Bilateral Settlement

Bilateral is the settlement that is being used in the current Internet. In *Bilateral*, the charging between two ISPs are determined by mutual agreements [7].

Let us first consider the charging between end-user and leaf tier T ISPs, which is the same in both settlement schemes. We assume that each end-user generates the same amount of traffic to the network, and is charged a fixed amount P_f by its provider ISP at tier T. The total amount of revenue collected from all end-users in the whole network is denoted as P_a.

For settlements between ISPs, we consider a per-traffic tier-based charging scheme. The precise charging scheme depends on the type of link between the two ISPs. Note that for simplicity we assume that the price is tier-dependent but not node-dependent. This is justified in a macroscopic view that ISPs in the same tier have similar size and thus have similar negotiation power in coming to a bilateral price. If the link is either a transit link or multi-homing link between tier t provider ISP and its tier $t+1$ customers, a charge P_t is paid by the customer to the provider for per-unit of traffic regardless of the traffic direction. If the link is a peering link between two ISPs, no charging is involved for any traffic sent over this link in either direction.

For the ease of later discussions, let us calculate the revenue obtained by an ISP n at tier t for handling (i.e., receiving and sending) one unit of traffic as the following:

$$\gamma_{n,t}(\boldsymbol{P}) = \begin{cases} 2P_t, & \text{for local traffic,} \\ P_t, & \text{for peering traffic,} \\ P_t - P_{t-1}, & \text{for upstream traffic,} \end{cases}$$

for $t \in [1, \ldots, T-1]$, and

$$\gamma_{n,T}(\boldsymbol{P}) = \begin{cases} -P_{T-1}, & \text{for upstream traffic,} \\ 0, & \text{for local/peering traffic,} \end{cases}$$

where $\boldsymbol{P} = \{P_t, \forall t \in \{1, \ldots, T-1\}\}$.

In this settlement method, each ISP keeps a traffic counter at its ingress links to record the total amount of traffic handled at each link. At a specific clearance time, typically once a month, ISP t charges their customer ISPs P_t times the traffic passed through their transit/multi-homing links.

B. Shapley Settlement

Shapley is the settlement that divides the revenue among ISPs according to Shapley value. It is previously studied in [2], where the authors showed several nice properties of implementing such settlement in the Internet.

Shapley value uses an axiomatic approach to allocate benefits obtained from a coalition among all participating players. Detailed information on Shapley value can be found in Shapley[8] and Winter [9].

Shapley is similar as *Bilateral* in terms of how the end-user is charged (i.e., flat-rate charging with P_f for each user and P_a for all users). The key difference is how different ISPs settle the charge among each other. Instead of charging differently for provider-customer relationship and peer-to-peer relationship, *Shapley* divides the revenue obtained from the end-users among ISPs by calculating the Shapley value of each ISP for each flow of traffic. The calculation of Shapley value is formally defined as follows:

Definition 1 (Shapley value): The Shapley value that node $n \in \boldsymbol{\mathcal{N}}$ obtains through handling traffic (a, b) (i.e., end-user a to end-user b) is:

$$\varsigma_n(a, b) = \frac{1}{|\boldsymbol{\mathcal{N}}|!} \sum_{\pi \in \boldsymbol{\mathcal{N}}} [v(\mathcal{S} \cup \{n\}, a, b) - v(\mathcal{S}, a, b)],$$

where $|\boldsymbol{\mathcal{N}}|$ is the total number of nodes in the network, π is the set of possible permutations of $\boldsymbol{\mathcal{N}}$, $\mathcal{S}$ is the subset of ISP nodes in a permutation that appears earlier than n, and $v(\mathcal{S}, a, b)$ is the characteristic function of Shapley value for the traffic from a to b under the sub-topology formed by the set $\mathcal{S}$. We first note that Shapley value is always between 0 and 1. Furthermore, Shapley value calculates the contribution of a player (a node in our case) as the normalized marginal contribution under all possible ways of including it into a set of nodes $\boldsymbol{\mathcal{N}}$. The measurement of worthiness is reflected by the characteristic function as defined next:

Definition 2 (Characteristic function of Shapley value): The characteristics function $v(.)$ for a set of nodes $\mathcal{S}$ handling a traffic from a to b is:

$$v(\mathcal{S}, a, b) = \begin{cases} 1 & \text{, if there exists a path in } \mathcal{S} \text{ that} \\ & \quad \text{connects } a \text{ to } b, \\ 0 & \text{, otherwise.} \end{cases}$$

This characteristic function concerns whether the traffic can be completed with set $\mathcal{S}$. When there is already a complete route from source to destination ($v(\mathcal{S}, a, b) = 1$), all extra nodes that join $\mathcal{S}$ later have zero marginal contribution ($v(\mathcal{S}', a, b) - v(\mathcal{S}, a, b) = 0, \mathcal{S}' \supset \mathcal{S}$), as they are considered redundant in terms of providing connectivity. Two examples of Shapley value is shown at Figure 2a and 2b.

To precisely implement the Shapley settlement, it is necessary to have a centralized network entity who calculates $\varsigma_n(a, b)$ for each ISP n for each traffic (a, b). At a specific clearance time, each ISP node uses the aggregated Shapley value calculated at the central entity to claim its portion of the total revenue P_a.

Based on the hierarchical model presented in this paper and BGP routing, it is possible to calculate the Shapley value based on a smaller set of parameters related to the traffic (a, b) as follows:

Definition 3: The Shapley value of n can be reformulated as

$$\psi_n(h, q, \boldsymbol{d}) = \varsigma_n(a, b),$$

where h is the height for traffic (a, b) (i.e. how far upstream the traffic goes), $q \in \{0, 1\}$ is the topmost peering parameter which states whether the traffic goes through a peering link at

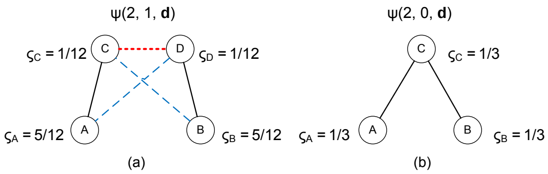

Fig. 2. An example of how Shapley value is calculated. We consider a traffic from an end-user of node A to an end-user of node B, and calculate the Shapley value of each node with respect to such traffic. Here the multi-homing factors are set as $\boldsymbol{d} : \{d_A, d_B\} = \{2, 2\}$ in (a) and $\{1, 1\}$ in (b).

its maximum height, and $\boldsymbol{d}$ is the set of multi-homing factor for all nodes $n \in \boldsymbol{\mathcal{N}}$.

Figure 2a and 2b show examples of how the parameters h and q are set for different traffic and topologies.

IV. Fair Price Achieving the Shapley Value: The Symmetric Case

The Shapley value is known as the fair way to distribute contributions among a group of players. However, adopting Shapley value in the current Internet is challenging since it requires the ISPs to first form a coalition, and besides the Shapley settlement requires centralized computation. Nevertheless, before we contemplate the likelihood of such a development, it would be helpful to find out the possibility of achieving the same revenue distribution as *Shapley* through the current Bilateral scheme, under suitably chosen prices.

In this section, we calculate the traffic-based price for each tier such that *Bilateral* produces the same revenue distribution as *Shapley*. We define the set of corresponding prices $\boldsymbol{P}^*$ as the *fair prices*:

Definition 4 (Fair price): The fair price $\boldsymbol{P}^* = \{P_t^*, t \in [1, \ldots, T-1]\}$ is the set of tier-based prices that make the expected revenue distribution in *Bilateral* equal to *Shapley* for all tiers T, i.e.,

$$P_a \sum_{n\in\mathcal{N}_t} \sum_{(a,b)\in\mathbf{R}} \mathbb{P}(a,b)\varsigma_n(a,b)$$
$$= \sum_{n\in\mathcal{N}_t} \sum_{(a,b)\in\mathbf{R}} \mathbb{P}(a,b)\gamma_{n,t}(\boldsymbol{P}^*) + P_a \times \mathbf{1}_{\{t=T\}}, \forall t \in \{1,\ldots,T\},$$

where $\mathbf{R}$ is the set of all possible source-destination pairs, $\mathbb{P}(a,b)$ is the probability that traffic (a,b) happens, and $\mathbf{1}_{\{.\}}$ is the indicator function.

In the rest of the paper, we will consider how the fair prices can be calculated in both symmetric and asymmetric topologies, as well as some nice properties of the fair prices (using symmetric topology as an example). After that, we will consider how the fair prices can be approximated in a distributed fashion based on only local information.

We first consider a symmetric topology, $\mathcal{G}^s = (T, \boldsymbol{\mathcal{N}}, \boldsymbol{\rho}^s, d)$, where $\boldsymbol{\rho}^s = \{\rho_t, t \in T\}$. In this topology, each node in the same tier t shares the same peering parameter ρ_t and all nodes of all tiers share the same multi-homing parameter d. For traffic parameters, we assume all ISPs have the same traffic parameter α, the total traffic intensity in the network is normalized to 1.

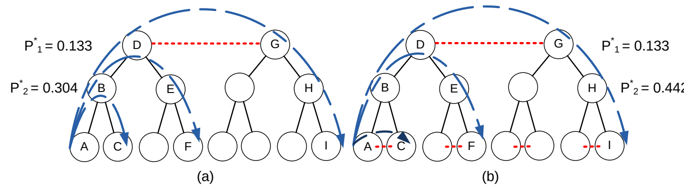

Fig. 3. Numerical examples of calculating fair price. We consider a three-tier model, the end-users are omitted in the figure. We let $P_a = 1$, $\alpha = 0.5$. Without loss of generality we draw all possible traffic (in terms of height travelled) initialized by node A by dash arrows. (a) shows the fair prices when $\rho_3 = 0$ and (b) show the fair prices when $\rho_3 = 1$.

Theorem 1: The fair price $\boldsymbol{P}^*(\mathcal{G}^s)$ in a symmetric hierarchical topology $\mathcal{G}^s = (T, \boldsymbol{\mathcal{N}}, \boldsymbol{\rho}^s, d)$ is:

$$P_t^*(\mathcal{G}^s) = \frac{P_a}{2(1-\alpha\rho_{t+1})}\left\{\Gamma_{\mathcal{G}^s}^{\text{tier1}} + \Gamma_{\mathcal{G}^s}^{\text{no-peer}} + \Gamma_{\mathcal{G}^s}^{\text{peer}}\right\}$$

$\forall t \in \{1, \ldots, T-1\}$,where

$$\Gamma_{\mathcal{G}^s}^{\text{tier1}} = 2t(1-\alpha)^t \psi_t(T, 1, d),$$
$$\Gamma_{\mathcal{G}^s}^{\text{no-peer}} = \sum_{i=1}^{t}(2i-1)\alpha(1-\alpha)^{i-1}(1-\rho_{t+2-i}) \times \psi_t(T-t+i, 0, d),$$
$$\Gamma_{\mathcal{G}^s}^{\text{peer}} = \sum_{i=1}^{t-1} 2i\alpha(1-\alpha)^i \rho_{t+1-i}\ \psi_t(T-t+i, 1, d).$$

This theorem is proofed by dynamic programming, which is omitted in this paper. Here $\Gamma_{\mathcal{G}^s}^{\text{tier1}}$ is the Shapley value corresponding to the traffic passing through the peered links in mesh at tier 1 times the probability of these traffic happens. $\Gamma_{\mathcal{G}^s}^{\text{no-peer}}$ is related to the Shapley value of all traffic going higher than tier t but do not use peering links at the highest height. $\Gamma_{\mathcal{G}^s}^{\text{peer}}$ is related to the Shapley value of traffic that passes through peering links at the traffic's highest height.

Through rearranging the equations in Theorem (1), we obtain the following relationship that sheds more light on the economic meaning of the fair price $\boldsymbol{P}^*$:

$$2(1-\alpha\rho_{t+1})P_t^* =$$
$$P_a\Big\{\sum_{(a,b)\in\mathbf{R}} \mathbb{P}(a,b) \sum_{n\in N_t} \varsigma_n(a,b) + 2(1-\alpha)(1-\alpha\rho_t)P_{t-1}^*\Big\},$$

By this formulation, the fair price at tier t is to account the contribution of tier t ISPs (first term on the right), plus the price that contains the contribution of the upper tiers (second term on the right) so that the upper tiers can achieve *Shapley* by charging tier t ISPs P_{t-1}^*. We then obtain the per-traffic based fair price by dividing the contributions by the total traffic intensity between all ISPs at tier t and tier $t+1$ (the term on the left).

Figure 3 shows a numerical example of calculating the fair prices. When peering links are established at tier 3 in Figure 3b, the fair price of tier 2 increases. When $\rho_3 = 1$, shorter routes like A to C will not pass through tier 2 ISP B. Since the fair price is the average per-traffic price of handling a traffic, losing the relatively cheap short distance traffic raises the average per-traffic price. Note that the actual revenue of tier 2 has decreased when $\rho_3 = 1$ as they handle less traffic.

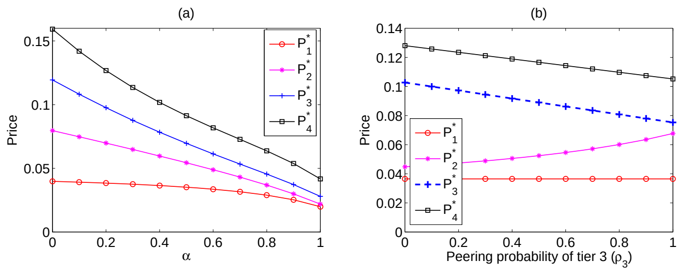

Fig. 4. (a) Fair price P_t^* verses α for different tier t: the fair price is decreasing in α. (b) P_t^* verses peering probability of tier 3 (ρ_3) for different tier t: the fair price of tier t is decreasing in ρ_t

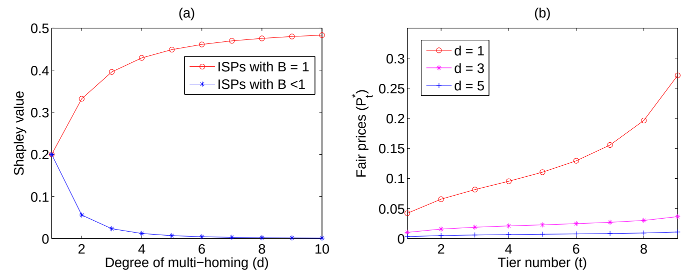

Fig. 5. (a) Shapley value ς for different multi-homing factor d for vital and non-vital ISPs: When d increases, vital nodes raises its contribution and non-vital nodes drops its contribution. (b) $\boldsymbol{P}^*$ for different multi-homing factor d: The fair price drops drastically when d increases.

V. Properties of the Fair Prices in the Symmetric Case

Under fair prices, *Bilateral* achieves the same revenue distribution as *Shapley*. Next we show that fair prices exhibits nice properties in the system parameters.

A. Sensitivity to traffic pattern α

Proposition 1: The fair price $\boldsymbol{P}^*$ is decreasing in $\alpha \in (0, 1]$ for any $t \in [1, \ldots, T-1]$.

The proposition is proofed by mathematical induction on the derivative of the fair price on α, which is not shown in this paper. Figure 4a shows the relationship between the fair price $\boldsymbol{P}^*$ and α by calculating the fair prices across α. When α increases, the traffic is more likely to go local, and thus is more likely to be handled by tiers that are closer to the leaf tiers. As a result, ISPs of tiers higher on the hierarchy have less chance to contribute to the connectivity of the traffic, thus it makes sense for them to charge less. On the other hand, when α decreases, a traffic tends to travel further away. In the extreme case where α approaches 0, all traffic travel to the farthest distance (i.e., through peering links in tier 1), and thus all tiers contribute equally to the connectivity and have the same revenue. This is reflected by the equal price difference across different tiers ($P_1^* - P_2^* = P_2^* - P_3^*$, etc.).

B. Sensitivity to network topology parameters ρ and d

Proposition 2: The fair price $\boldsymbol{P}^*$ is decreasing in $\rho_t \in (0, 1)$ for any $t \in [1, \ldots, T-1]$.

The proof is derived by a simple differentiation that is not shown here. Figure 4b shows an mathematical example of how the fair prices of different tiers change as a function of the peering probability of tier 3 ISPs ρ_3. When ρ_3 increases, more traffic passes through the peering link in tier 3 without using services from their provider ISPs. Paying less to the upstream ISPs reduces the fair price of tier 3 ISPs. With the same argument, the fair price of the lower tiers below tier 3 reduces as well.

The fair price of tier 2, on the other hand, increases when ρ_3 increases. As explained in the numerical example in Figure 3, it is because the average traffic price increases when tier 2 ISPs handle less shorter distance traffic.

Before we mention the sensitivity of the multi-homing factor d, we could like to introduce the concept of *betweenness*, as an indicator to characterize the importance of an ISP node as follows:

Definition 5 (Betweenness): The betweenness of a node n with respect to traffic (a, b) is:

$$\mathcal{B}_n(a, b) = \frac{\sigma(a, n, b)}{\sigma(a, 0, b)},$$

where $\sigma(a, n, b)$ is the number of routes from a to b that pass through node n. If $n = 0$, then $\sigma(a, 0, b)$ denotes the total number of routes from a to b. It is clear that $\mathcal{B}_a(a, b)$ and $\mathcal{B}_b(a, b)$ are always equal to 1.

Observation 1: The Shapley value of an ISP n increases in multi-homing factor d if its betweenness $\mathcal{B}_n = 1$, and decreases in d if its betweenness $\mathcal{B}_n < 1$. Furthermore, $\boldsymbol{P}^* \to 0$ when $d \to \infty$.

Figure 5a shows the differences in Shapley value between a vital player (betweenness $\mathcal{B} = 1$) and a non-vital player ($\mathcal{B} < 1$) for varies d. When d increases, the Shapley value of non-vital ISPs decreases significantly since there are more alternative ways for providing connectivity, while the Shapley value of vital ISPs increases significantly since they remain the necessary components for any choice of connectivity.

Figure 5b shows that $\boldsymbol{P}^*$ drops dramatically when d increases and flattens quickly for $d \geq 5$, as this connectivity-focused characteristic function gives little value to non-vital ISPs. For $d \geq 1$, the only vital players for each traffic are the initiating ISP and the terminating ISP. In this scenario, initiating leaf ISP and the terminating leaf ISP at tier T collect most of the revenue from handling the traffic, leaving the rest of the non-vital upper tier ISPs little revenue to share. This explains why we have $\boldsymbol{P}^* \to 0$ when $d \to \infty$. In this extreme case, there are infinite number of paths that can be used to transmit the same traffic, thus the contribution of a non-vital ISP on any particular path to connectivity is virtually zero.

VI. Fair Price Achieving the Shapley Value: The Asymmetric Case

Here we relax the symmetric network assumption in Sections V and IV, and show how the fair price is determined in the asymmetric case. Here we consider a general topology model $\mathcal{G} = (T, \boldsymbol{\mathcal{N}}, \boldsymbol{\rho}, \boldsymbol{d})$, where $\boldsymbol{\rho} \in \{\rho_n, \forall n \in \boldsymbol{\mathcal{N}}\}$ and $\boldsymbol{d} \in \{d_n, \forall n \in \boldsymbol{\mathcal{N}}\}$. In terms of traffic model, each ISP node n uses parameters (α_n, I_n) to describe the traffic passing through it. The *fair price* in this asymmetric case, $\boldsymbol{P}^*(\mathcal{G})$, can be calculated as follows:

Theorem 2: The fair price of an ISP at tier t in an asymmetric topology $\mathcal{G} = (T, \boldsymbol{\mathcal{N}}, \boldsymbol{\rho}, \boldsymbol{d})$ is:

$$P_t^*(\mathcal{G}) = \frac{P_a}{|\mathcal{N}_t|} \sum_{m \in \mathcal{N}_t} \sum_{n \in \mathcal{N}_{t+1}^m} \frac{I_n}{2(1-\alpha_n \rho_n)} \times \left\{ \Gamma_{\mathcal{G}}^{\text{tier1}} + \Gamma_{\mathcal{G}}^{\text{no-peer}} + \Gamma_{\mathcal{G}}^{\text{peer}} \right\} \forall t \in \{1, \ldots, T-1\},$$

where $\mathcal{N}_t^m$ is the set of ISPs at tier t that has ISP m as its provider ISP(through either transit links or multi-homing links). The three pricing terms in Theorem (2) can be similarly derived as in Theorem (1), and is not shown in this paper.

VII. Distributed and Local Approximation of the Fair Price

In order to calculate the fair prices based on either Theorem (1) (symmetric case) or Theorem (2) (asymmetric case), we require global information such as the full network topology, which is impractical to obtain in the Internet. Next we illustrate a distributed approximation scheme that allows each node to calculate the fair price with limited local information.

For a node n at tier t, it calculates an approximated fair price $\boldsymbol{P'}^*$ based on the following inputs: 1) the price charged by its provider ISPs at tier $t-1$ (P_{t-1}^*); 2) the incoming traffic pattern $(\alpha_m, I_m), \forall m \in \mathcal{N}_t$; 3) the betweenness $\mathcal{B}_m$ of all nodes $m \in \mathcal{N}_t$, and 4) a heuristic approximation of peering parameters $\rho'_t, t \in \{2, \ldots, t-1\}$.

The fair price can be approximated by

$$P'^*_t(\mathcal{G}) = \frac{P_a}{|\mathcal{N}_t|} \sum_{m \in \mathcal{N}_t} \sum_{n \in \mathcal{N}_{t+1}^m} \frac{I_n}{2(1-\alpha_n \rho_n)} \left\{ \Gamma_{\mathcal{G}}^{\prime\text{tier1}} + \Gamma_{\mathcal{G}}^{\prime\text{guess}} + \Gamma_{\mathcal{G}}^{\prime\text{tier}t} + \Gamma_{\mathcal{G}}^{\text{Bilateral}} \right\},$$

where $\Gamma_{\mathcal{G}}^{\prime\text{tier1}}$ and $\Gamma_{\mathcal{G}}^{\prime\text{tier}t}$ are the approximated Shapley value of traffic to tier 1 ISPs and tier t ISPs respectively. $\Gamma_{\mathcal{G}}^{\prime\text{guess}}$ is the approximated Shapley value of traffic that goes to heights between tier 2 and tier $t-1$, which requires ρ'_t for approximation. $\Gamma_{\mathcal{G}}^{\text{Bilateral}}$ is the precise price that tier t ISPs pay their provider ISPs for all upstream traffic.

This approximation is based on two observations. First, we approximate the original Shapley value function $\psi_n(T, q, \boldsymbol{d})$ by $\psi'_n(T, \mathcal{B}_n)$, which requires only local information to compute. The second observation is that a general idea of how each tier peers is feasible to obtain, and sufficient to calculate a close fair price. As shown in Figure 4b, the fair prices do not significantly change for two close ρ_t, therefore we can use a heuristic approximation to replace this piece of information.

To illustrate the performance of the approximation, we simulate a four-tier topology where the real peering parameter and heuristic approximation of the peering pattern are both independently and randomly generated. We use the following approximation for the Shapley value calculation:

$$\psi'_n(T, \mathcal{B}_n) = \frac{1}{2}\left[\psi_n(T, 1, \boldsymbol{d'}_n) + \psi_n(T, 0, \boldsymbol{d'}_n)\right],$$

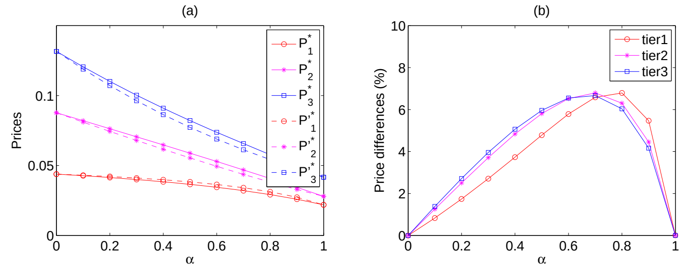

Fig. 6. Simulated performance of the distributed approximation scheme for the fair price. Here we consider a four-tier topology with randomly chosen true and approximated peering parameters. (a) the precise and approximated fair prices; (b) the difference between real and approximated fair prices of different tiers in percentage.

where $\boldsymbol{d'}_n = \{3, \ldots, d_n, \ldots, 3\}$, i.e., assuming all tiers have the same multi-homing factor equal to 3, except the value at the $t+1$ tier which is approximated by d_n.

Figure 6 shows that the approximation is very accurate in this random case, where the difference is no larger than 8%. If the peering parameters are estimated according to some network statistics instead of just random guessing, then the performance of the approximation is likely to be further improved.

VIII. Conclusion

In this paper, we try to ask the following question: is it possible to set the proper bilateral prices so that a fair revenue distribution among ISPs can be achieved using the revenue settlement mechanism used in the current Internet? Through both careful analysis and extensive simulations, we have obtained positive preliminary answers to this question. We show that by setting the bilateral charging based on global topology and traffic information for each ISP tier, the current bilateral settlement distribution leads to Shapley value, which is widely regarded as a fair cooperative revenue distribution method. We further show that such fair prices are reasonable to be adopted in the Internet, as it shows sensible properties to parameter changes. In terms of feasibility to be implemented, we show that the fair price can be approximated using local information.

References

[1] J. Musacchio, S. G., and W. J., "Network Neutrality and Provider Investment Incentives," *Proc. Network Economics*, 2007.

[2] R. Ma, D. Chiu, J. Lui, V. Misra, and D. Rubenstein, "Internet Economics: The use of Shapley value for ISP settlement," *Proc. CoNEXT*, 2007.

[3] G. Huston, "Interconnection, peering, and settlements," *Proc. INET*, 1999.

[4] R. Johari and J. Tsitsiklis, "Routing and peering in a competitive Internet," *Prof. IEEE CDC*, vol. 2, 2004.

[5] W. Norton, "Internet Service Providers and Peering," *Proceedings of NANOG*, 2001.

[6] L. Gao, "On inferring autonomous system relationships in the internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 733–745, 2001.

[7] J. Bailey, "The Economics of Internet Interconnection Agreements," *Internet Economics*, pp. 155–168, 1995.

[8] L. Shapley and A. Roth, *The Shapley Value: Essays in Honor of Lloyd S. Shapley*. Cambridge University Press, 1988.

[9] E. Winter, "The Shapley Value," *Handbook of Game Theory with Economic Applications*, vol. 3, pp. 2025–2054, 2002.