

On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers*

Richard T.B. Ma
Columbia University
tbma@ee.columbia.edu

Dah-ming Chiu
The Chinese University of HK
dmchiu@ie.cuhk.edu.hk

John C.S. Lui
The Chinese University of HK
cslui@cse.cuhk.edu.hk

Vishal Misra
Columbia University
misra@cs.columbia.edu

Dan Rubenstein
Columbia University
danr@ee.columbia.edu

ABSTRACT

Internet service providers (ISPs) depend on one another to provide global network services. However, the profit-seeking nature of the ISPs leads to selfish behaviors that result in inefficiencies and disputes in the network. This concern is at the heart of the "network neutrality" debate, which also asks for an appropriate compensation structure that satisfies all types of ISPs. Our previous work showed in a general network model that the Shapley value has several desirable properties, and that if applied as the revenue model, selfish ISPs would yield globally optimal routing and interconnecting decisions.

In this paper, we use a more detailed and realistic network model with three classes of ISPs: content, transit, and eyeball. This additional detail enables us to delve much deeper into the implications of a Shapley settlement mechanism. We derive closed-form Shapley values for more structured ISP topologies and develop a dynamic programming procedure to compute the Shapley revenues under more diverse Internet topologies. We also identify the implications on the bilateral compensation between ISPs and the pricing structures for differentiated services. In practice, these results provide guidelines for solving disputes between ISPs and for establishing regulatory protocols for differentiated services and the industry.

1. INTRODUCTION

*This research was supported in part by the National Science Foundation grants CNS-0627590, ITR-0325495 and CNS-0238299. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

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 2008, December 10-12, 2008, Madrid, SPAIN
Copyright 2008 ACM 978-1-60558-210-8/08/0012 ...\$5.00.

The Internet is operated by thousands of interconnected ISPs, with each ISP interested in maximizing its own profit. Rather than operating independently, each ISP requires the cooperation of other ISPs in order to provide Internet services. However, without an appropriate profit sharing mechanism, profit-seeking objectives often induce various selfish behaviors in routing [21] and interconnecting [6], degenerating the performance of the network. For example, Level 3 unilaterally terminated its "settlement free" peering relationship with Cogent on October 5, 2005. This disruption resulted in at least 15% of the Internet to be unreachable for the users who utilized either Level 3 or Cogent for Internet access. Although both companies restored peering connections several days later with a new on-going negotiation, Level 3's move against Cogent exhibited an escalation of the tension that necessitates a new settlement for ISPs.

Compared to the traditional settlement models [1, 9] in telecommunication, the Internet architecture has exhibited a more versatile and dynamic structure. The most prevalent settlements a decade ago were in the form of bilateral negotiations, with both parties creating either a *customer/provider* or a *zero-dollar peering* relationship [9]. Today, because of the heterogeneity in ISPs, simple peering agreements are not always satisfactory to all parties involved, and *paid peering* [5] has naturally emerged as the preferred form of settlement among the heterogeneous ISPs. Nevertheless, the questions like "which ISP should pay which ISP?" and "how much should ISPs pay each other?" are still unsolved. These open questions are also closely related to the *network neutrality* [3, 25, 7] debate, which argues the appropriateness of providing service/price differentiations in the Internet.

Our previous work [12] explored the application of *Shapley value* [24, 20], a well-known economic concept originated from coalition games [19, 4, 10], to a general network setting. We proved that if profits were shared as prescribed by the Shapley value mechanism, not only would the set of desirable properties inherent to the Shapley solution exist, but also that the selfish behaviors of the ISPs would yield globally optimal routing and interconnecting decisions.

In this paper, we explore the Shapley value revenue distribution in a detailed Internet model and its implications on

the stability of prevalent bilateral settlements and the pricing structure for differentiated services in the Internet. Faratin et al. [5] view today's Internet as containing two classes of ISPs: *content* and *eyeball*. Content ISPs specialize in providing hosting and network access for end-customers and commercial companies that offer content, such as Google, Yahoo, and YouTube. Eyeball ISPs, such as AT&T and Verizon, specialize in delivery to hundreds of thousands of end-customers, i.e., supporting the last-mile connectivity. Our previous work [13] explored the Shapley value revenue distribution based on this Content-Eyeball (CE) model. This paper significantly extends and generalizes the results from the CE model (Theorem 1). We start with the CE model and extend it to include a third class of ISPs: transit ISPs. Transit ISPs model the Tier-1 ISPs, such as Level 3 and Cogent, which provide transit services for other ISPs and naturally form a full-mesh topology to provide the universal accessibility of the Internet. Our new results are:

- We obtain closed-form Shapley value solutions for ISPs in the Content-Transit-Eyeball (CTE) models under bipartite topologies (Theorem 2).
- We generalize the closed-form Shapley solution for multiple contents/regions environments where “inelastic” components can be decomposed linearly (Theorem 3¹).
- We derive a dynamic programming procedure to calculate the Shapley value for ISPs under general Internet topologies. This procedure can progressively build up the Shapley values for ISPs along with the development of the network structure (Theorem 4).
- We show that the aggregate revenue can be decomposed by content-side payments and eyeball-side payments, each of which can be distributed as a Shapley value revenue of a *canonical* subsystem to each ISP that contributed in the coalition.
- Through the Shapley value solution, we explain 1) why the zero-dollar peering and customer/provider bilateral agreements can be stable in the early stage of the Internet, 2) why paid-peering has emerged recently, and 3) why an unconventional reverse customer/provider relationship should exist in order for the bilateral agreements to be stable.
- Instead of supporting or disproving service differentiations in the network neutrality debate, we try to answer the question what the appropriate pricing structure is for differentiated services that are proven to be beneficial to the society. Based on the Shapley value solution, we discuss the implied compensation structures for potential applications of differentiated services.

We believe that these results provide guidelines for ISPs to settle bilateral disputes, for regulatory institutions to design

¹All results in [13] can be considered as special cases of a multiple contents/regions environment with only content and eyeball ISPs.

pricing regulations, and for developers to negotiate and cooperate to provide differentiated services on top of the current Internet.

2. SHAPLEY VALUE AND PROPERTIES

Here, we briefly introduce the concept of Shapley value and its use under our ISP revenue distribution context. We follow the notations in [12]. We consider a network system comprised of a set of ISPs denoted as \mathcal{N} . $N = |\mathcal{N}|$ denotes the number of ISPs in the network. We call any nonempty subset $\mathcal{S} \subseteq \mathcal{N}$ a *coalition* of the ISPs. Each coalition can be thought of as a sub-network that might be able to provide partial services to their customers. We denote v as the *worth function*, which measures the monetary payments produced by the sub-networks formed by all coalitions. In other words, for any coalition \mathcal{S} , $v(\mathcal{S})$ defines the revenue generated by the sub-network formed by the set of ISPs \mathcal{S} . In particular, v measures the aggregate payments obtained by ISPs in a coalition as

$$v(\mathcal{S}) = \sum_{i \in \mathcal{S}} P_i(\mathcal{S}), \quad (1)$$

where $P_i(\mathcal{S})$ is the end-payment collected by ISP i in coalition \mathcal{S} . Thus, the network system is defined as the pair (\mathcal{N}, v) . Through the worth function v , we can measure the contribution of an ISP to a group of ISPs as the following.

Definition 1. *The marginal contribution of ISP i to a coalition $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$ is defined as $\Delta_i(v, \mathcal{S}) = v(\mathcal{S} \cup \{i\}) - v(\mathcal{S})$. Proposed by Lloyd Shapley [24, 20], the Shapley value serves as an appropriate mechanism for ISPs to share revenues.*

Definition 2. *The Shapley value φ is defined by*

$$\varphi_i(\mathcal{N}, v) = \frac{1}{N!} \sum_{\pi \in \Pi} \Delta_i(v, S(\pi, i)) \quad \forall i \in \mathcal{N}, \quad (2)$$

where Π is the set of all $N!$ orderings of \mathcal{N} and $S(\pi, i)$ is the set of players preceding i in the ordering π .

The Shapley value of an ISP i can be interpreted as the expected marginal contribution $\Delta_i(v, \mathcal{S})$ where \mathcal{S} is the set of ISPs preceding i in a uniformly distributed random ordering. The Shapley value depends only on the values $\{v(\mathcal{S}) : \mathcal{S} \subseteq \mathcal{N}\}$, and satisfies a bunch of desirable efficiency and fairness properties [12].

We showed in [12] that the Shapley value mechanism also induces Nash equilibria that are globally optimal for routing and interconnecting. However, the calculation of the Shapley value involves an exponential time complexity. In this paper, we focus on the calculation of the Shapley value revenues for ISPs and the implications derived from the Shapley solution. We assume that routing costs are negligible compared to the revenue obtained from providing services. Nevertheless, our framework can always be extended to include an orthogonal direction of routing decisions and costs.

3. NETWORK MODEL

Faratin et al. [5] categorize ISPs as two basic types: content ISPs and eyeball ISPs. We extend this categorization by including a third type: transit ISPs. The set of ISPs is defined as $\mathcal{N} = \mathcal{C} \cup \mathcal{T} \cup \mathcal{B}$, where $\mathcal{C} = \{C_1, \dots, C_{|\mathcal{C}|}\}$ denotes the set of content ISPs, $\mathcal{T} = \{T_1, \dots, T_{|\mathcal{T}|}\}$ denotes the set of transit ISPs, and $\mathcal{B} = \{B_1, \dots, B_{|\mathcal{B}|}\}$ denotes the set of eyeball ISPs. We denote \mathcal{Q} as the set of contents provided by the set of content ISP \mathcal{C} . Each content ISP C_i provides a subset of the contents $Q_i \subseteq \mathcal{Q}$. The intersection of any Q_i and $Q_{i'}$ might not be empty, meaning C_i and $C_{i'}$ can provide duplicate contents. We denote \mathcal{R} as the set of regions covered by the set of eyeball ISP \mathcal{B} to provide Internet services to end-customers. Each eyeball ISP B_j covers a subset of the regions $R_j \subseteq \mathcal{R}$. We assume that each region $r \in \mathcal{R}$ has a fixed end-customer population of size X_r . Each end-customer chooses one of the eyeball ISPs covering the region for service; therefore, each eyeball ISP B_j attracts and serves a portion x_j^r (equals zero if B_j does not cover region r , i.e. $r \notin R_j$) of the total population in region r . We assume that each content or eyeball ISP is connected to one (single-homing) or multiple (multi-homing) transit ISPs; while, transit ISPs connect with one another, forming a full-mesh topology. We denote CP_i as the content-side payment from content providers to ISP C_i and BP_j as the eyeball-side payment from end-customers to ISP B_j .

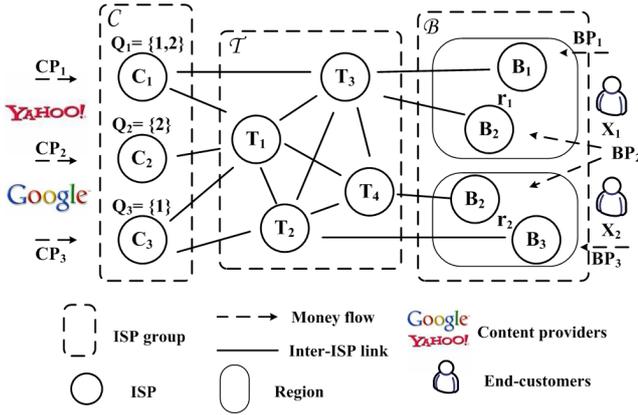


Figure 1: The Content-Transit-Eyeball ISP model.

Figure 1 illustrates a scenario with $|\mathcal{C}| = |\mathcal{B}| = 3$, $|\mathcal{T}| = 4$, $\mathcal{Q} = \{1, 2\}$, and $\mathcal{R} = \{1, 2\}$. The contents provided by three content ISPs are $Q_1 = \{1, 2\}$, $Q_2 = \{2\}$ and $Q_3 = \{1\}$ respectively. Eyeball B_1 and B_3 covers region R_1 and R_2 respectively, i.e. $R_1 = \{1\}$, $R_3 = \{2\}$, and B_2 covers both regions, i.e. $R_2 = \{1, 2\}$.

The justification of our network model to represent the Internet ISP structure is supported by the study from the Cooperative Association for Internet Data Analysis (CAIDA) [16]. Their study shows that the average distance of AS-level topology is less than 4, and 62% of AS paths are 3-hop paths. This suggests that the most frequent path patterns are described in our Content-Transit-Eyeball model: data source originates from content ISPs, going through either one or two transit ISPs and reaches eyeball ISPs. The full-mesh

topology of transit ISPs is also true for Tier-1 ISPs to provide universal accessibility of the Internet in reality [5].

3.1 Revenue Model

We define each eyeball-side payment BP_j as follows.

$$BP_j = \sum_{r \in R_j} \alpha_r x_j^r \quad \forall B_j \in \mathcal{B}, \quad (3)$$

where α_r is the monthly charge in region r . The eyeball-side payment is basically the aggregate service charges from end-customers served by ISP B_j in different regions. We assume that the monthly charge might be different in distinct regions due to various economic and living conditions; however, within the same region, we assume that the market is competitive so that eyeball ISPs charge the same price for users. We also define each content-side payment CP_i as follows.

$$CP_i = \sum_{q \in Q_i} \beta_q \rho_i^q \quad \forall C_i \in \mathcal{C}, \quad (4)$$

where β_q is the average per customer revenue generated by uploading content q , and ρ_i^q is the aggregate amount of end-customers that might request content q from content ISP C_i . Content-side payments are paid by content-providers who run business through the Internet. Content ISPs receive payment by disseminating business information, e.g. product information and advertisements, for content-providers. In practice, β_q relates to two factors. First, β_q measures the popularity of content q . When the content is popular, larger percentage (higher download frequency) of the customer size ρ_i^q will download it, generating more revenue for the content ISPs. Second, β_q measures the revenue for delivering per-unit information of content q . The higher the importance of the content to the provider, the higher the revenue paid by them. For example, Ebay or Amazon might pay more for products information browsed by customers, which potentially generates profits for them.

The above BP_j and CP_i are defined on general customer distribution $\{x_j^r : r \in \mathcal{R}, j \in \mathcal{B}\}$ and content demand distribution $\{\rho_i^q : q \in \mathcal{Q}, i \in \mathcal{C}\}$. Before describing both distributions, we first define two similar variables as follows:

$$h_q = \sum_{i=1}^{|\mathcal{C}|} \mathbf{1}_{\{q \in Q_i\}} \quad \forall q \in \mathcal{Q}, \quad \text{and} \quad h_r = \sum_{j=1}^{|\mathcal{B}|} \mathbf{1}_{\{r \in R_j\}} \quad \forall r \in \mathcal{R}.$$

These two variables define the number of content ISPs that provide content q and the number of eyeball ISPs that cover region r respectively. Notice that with ISP arrival or departure, these variables change accordingly.

3.2 Demand Assumptions

We assume that the inter-region customer demand is *inelastic* and the intra-region customer demand is *elastic*. By “inelastic”, we mean the total demand in each region r is the fixed population X_r , regardless the number of eyeball ISPs covering the region (as long as $h_r > 0$). In practice, it models the situation that the customers in a certain region

cannot physically choose ISPs from other regions, even if the choices of local ISPs are limited. However, by “elastic”, we mean that customers are not sticky to a certain eyeball ISP within a region. When a new eyeball ISP comes into a region, h_r increases by one and customers from the original ISPs may shift to the new ISP. Similarly, when an eyeball ISP leaves a region, its customers will choose the remaining ISPs evenly. In both cases, the total customer size of a region is fixed.

We also assume that any end-customer will download a particular content q from any of content ISPs that provide the content. Conceptually, we can imagine that the same content provided by multiple content ISPs as an “elastic” content where users can choose which ISP to download from. Similarly, different content can be thought of as “inelastic” where users have to download a certain content specifically from the group of content ISPs that provide it.

3.3 Conservation of Revenue

In general, the total amount of revenue generated by the whole network is a constant:

$$v(\mathcal{N}) = \sum_{j=1}^{|\mathcal{B}|} BP_j + \sum_{i=1}^{|\mathcal{C}|} CP_i = \sum_{r \in \mathcal{R}} \left(\alpha_r + \sum_{q \in \mathcal{Q}} \beta_q \right) X_r. \quad (5)$$

However, the whole network might have revenue loss if the network is segmented so that some customer cannot reach all content provided by all content ISPs. Figure 2 illus-

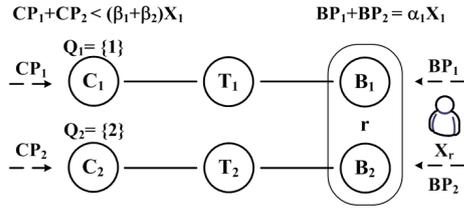


Figure 2: A segmented network results revenue loss.

trates an example where customers, attached to one of the eyeball ISPs, can only access one of the content ISPs. The content-side revenue will be less than $\sum_{q \in \mathcal{Q}} \beta_q (\sum_{r \in \mathcal{R}} X_r)$. However, if the two transit ISPs are connected, the total revenue will follow Equation (5). In this paper, we assume the condition of Equation (5) holds whenever there is topological changes in the network (i.e. interconnection change and ISP arrival/departure). The justification of this assumption is that, in practice, the set of transit ISPs are made up of Tier-I ISPs that always connect to each other forming a full mesh [5, 16]. Therefore, $v(\mathcal{N})$ really defines the total revenue generated by the Internet.

4. SHAPLEY REVENUE DISTRIBUTION

In this section, we progressively develop the Shapley value revenue distribution for ISPs under different models. We start with a single-content/single-region scenario with a well-connected topology. We first consider a model with only

content and eyeball ISPs. After that, we extend the result to include transit ISPs. Then, we extend the result for a multiple-content/multiple-region model. Finally, we explore more general Internet topologies with previous models.

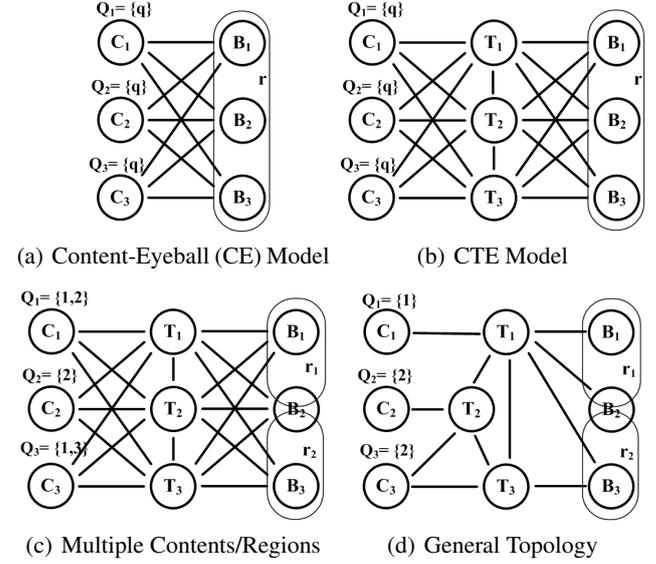


Figure 3: Progressively developed models.

Figure 3 illustrates the four models we are going to discuss in the following subsections.

4.1 Content-Eyeball (CE) Model

The Content-Eyeball (CE) model follows from Faratin et al. [5]. We focus on single-content/single-region for the time being, assuming that every content ISP provides the same content q and every eyeball ISP covers the same region r . As illustrated in Figure 3(a), we also focus on the topology where content and eyeball ISPs form a complete bipartite graph. We have $|\mathcal{N}| = |\mathcal{C}| + |\mathcal{B}|$ and $v(\mathcal{N}) = (\alpha_r + \beta_q)X_r$. We define φ_{B_j} and φ_{C_i} as the Shapley value revenue distributed to B_j and C_i respectively, and $\phi_B = \sum_{B_j \in \mathcal{B}} \varphi_{B_j}$ and $\phi_C = \sum_{C_i \in \mathcal{C}} \varphi_{C_i}$ as the aggregate Shapley value for the group of eyeball and content ISPs respectively. Due to the space limitation, the proofs of theorems are omitted but are available in the technical report [14].

Theorem 1 (SHAPLEY VALUE FOR CE MODEL). *We consider a set of content ISPs, \mathcal{C} , providing one content and a set of eyeball ISPs, \mathcal{B} , covering one region. Under the CE model with a complete bipartite graph topology, the Shapley value revenues for each ISP are the following:*

$$\varphi_{B_j}(|\mathcal{B}|, |\mathcal{C}|) = \frac{|\mathcal{C}|}{|\mathcal{B}||\mathcal{N}|} v(\mathcal{N}) \quad \forall B_j \in \mathcal{B},$$

$$\varphi_{C_i}(|\mathcal{B}|, |\mathcal{C}|) = \frac{|\mathcal{B}|}{|\mathcal{C}||\mathcal{N}|} v(\mathcal{N}) \quad \forall C_i \in \mathcal{C}.$$

Theorem 1 shows that an ISP’s Shapley value is inversely proportional to the number of ISPs of the same type, and

proportional to the number of ISPs of the opposite type. In particular, the aggregate Shapley value revenue of both types of ISPs are inverse proportional to the number of ISPs of each type, i.e. $\phi_C : \phi_B = |\mathcal{B}| : |\mathcal{C}|$.

Corollary 1 (MARGINAL REVENUE). *Suppose any content ISP de-peers with all eyeball ISPs, i.e. removing some $C_i \in \mathcal{C}$ from \mathcal{C} . We define $\mathcal{C}' = \mathcal{C} \setminus \{C_i\}$ as the set of remaining content ISPs. The marginal aggregate revenue for the set of eyeball ISPs is defined as*

$$\Delta_{\phi_B} = \phi_B(|\mathcal{B}|, |\mathcal{C}'|) - \phi_B(|\mathcal{B}|, |\mathcal{C}|).$$

This marginal aggregate revenue satisfies:

$$\Delta_{\phi_B} = -\frac{|\mathcal{B}|\phi_B}{|\mathcal{C}|(|\mathcal{N}| - 1)} = -\frac{|\mathcal{B}|v(\mathcal{N})}{|\mathcal{N}|(|\mathcal{N}| - 1)}.$$

Corollary 1 measures the marginal revenue loss for the set of eyeball ISPs from losing one of content ISPs. Because the Shapley revenue in Theorem 1 is symmetric between content and eyeball ISPs, we can derive similar marginal revenue result by considering any de-peering of an eyeball ISP as well. When $|\mathcal{C}| = 1$, Corollary 1 tells that the marginal revenue for the group of eyeball ISPs is $-\phi_B$, which indicates that if the only content ISP leaves the system, all eyeball ISPs are going to lose all their revenue. When $|\mathcal{B}| = 1$, Corollary 1 tells that $\Delta_{\phi_B} = -\frac{1}{|\mathcal{C}|^2}\phi_B$, which mean that by disconnecting to an additional content ISP, this eyeball ISP is going to lose $1/|\mathcal{C}|^2$ of its original revenue. The CE model gives a good sense of the Shapley value revenue among ISPs before the model is getting more complicated. However, we will see that more detailed models show similar revenue-sharing features as the basic model.

4.2 Content-Transit-Eyeball (CTE) Model

The Content-Transit-Eyeball (CTE) model, as illustrated in Figure 3(b), extend the CE model by introducing a set of transit ISPs in between the content and eyeball ISPs. Again, the topology between two connecting classes of ISPs are assumed to be a complete bipartite graph. Although the transit ISPs are supposed to form a full-mesh, the CTE model does not put any constraint on the interconnections between any pair of transit ISPs. Because any content ISP can be reached by any eyeball ISP via exactly one of the transit ISPs, interconnecting links between transit ISPs are “dummy” in this topology in that, their presence does not affect the Shapley revenue of any ISP. Later, we will extend our model to general Internet topologies where transit ISPs do require to form a full mesh and eyeball ISPs might need to go through multiple transit ISPs to reach certain content ISPs. Here, the total number of ISPs in the network becomes $|\mathcal{N}| = |\mathcal{C}| + |\mathcal{T}| + |\mathcal{B}|$. Similarly, we define φ_{T_k} as the Shapley value revenue of T_k , and $\phi_{\mathcal{T}} = \sum_{T_k \in \mathcal{T}} \varphi_{T_k}$ as the aggregate Shapley value for the group of transit ISPs.

Theorem 2 (SHAPLEY VALUE FOR CTE MODEL). *We consider a network with a set of content ISPs \mathcal{C} , a set of eyeball ISPs \mathcal{B} and a set of transit ISPs \mathcal{T} . Both content and*

eyeball ISPs are connected to the transit ISPs by a complete bipartite graph. Assume all content ISPs provide a single content and all eyeball ISPs cover a single region. The Shapley value revenue for each ISP is in the following form:

$$\varphi_{B_j}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) = \varphi_B(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|)v(\mathcal{N}) \quad \forall B_j \in \mathcal{B},$$

$$\varphi_{T_k}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) = \varphi_T(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|)v(\mathcal{N}) \quad \forall T_k \in \mathcal{T},$$

$$\varphi_{C_i}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) = \varphi_C(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|)v(\mathcal{N}) \quad \forall C_i \in \mathcal{C},$$

where the normalized Shapley values φ_B, φ_T and φ_C are:

$$\varphi_B(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) = \frac{1}{|\mathcal{N}|} \sum_{t=1}^{|\mathcal{T}|} \sum_{c=1}^{|\mathcal{C}|} \binom{|\mathcal{T}|}{t} \binom{|\mathcal{C}|}{c} \binom{|\mathcal{N}| - 1}{t+c}^{-1},$$

$$\varphi_T(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) = \frac{1}{|\mathcal{N}|} \sum_{b=1}^{|\mathcal{B}|} \sum_{c=1}^{|\mathcal{C}|} \binom{|\mathcal{B}|}{b} \binom{|\mathcal{C}|}{c} \binom{|\mathcal{N}| - 1}{b+c}^{-1},$$

$$\varphi_C(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) = \frac{1}{|\mathcal{N}|} \sum_{b=1}^{|\mathcal{B}|} \sum_{t=1}^{|\mathcal{T}|} \binom{|\mathcal{B}|}{b} \binom{|\mathcal{T}|}{t} \binom{|\mathcal{N}| - 1}{b+t}^{-1}.$$

The normalized Shapley values φ_B, φ_T and φ_C in Theorem 2 can be considered as the percentage share of the total revenue $v(\mathcal{N})$ for each ISP. Theorem 2 shows that φ_B, φ_T and φ_C are symmetric (which is also true for the CE model), in the sense that they can be represented by the same function with arguments shuffled:

$$\begin{aligned} \varphi_B(b, t, c) &= \frac{1}{|\mathcal{N}|} \sum_{t'=1}^t \sum_{c'=1}^c \binom{t}{t'} \binom{c}{c'} \binom{b-1+t+c}{t'+c'}^{-1} \\ &= \varphi_T(t, b, c) = \varphi_C(c, t, b). \end{aligned}$$

To understand this symmetric property of the normalized Shapley values function, we can think the importance of the group of transit ISPs is the same as that of either the group of content ISPs or the eyeball ISPs, because without the transit ISPs, the network is totally disconnected and cannot generate any revenue.

Figure 4 plots the aggregate Shapley value revenue of the set of transit ISPs, $\phi_{\mathcal{T}}$, against different sizes of the eyeball and the content ISPs. We normalize $v(\mathcal{N})$ to be 1. On the x-axis, we vary the size of the content ISPs, $|\mathcal{C}|$. On each plotted curve, the size of the transit ISPs $|\mathcal{T}|$ is a constant. With the change of $|\mathcal{C}|$, the size of the eyeball ISPs changes accordingly to satisfy the relation $|\mathcal{B}| = |\mathcal{T}| - |\mathcal{C}|$. Effectively, when we increase the number of content ISPs, we decrease the number of eyeball ISPs and keep the number of transit ISPs the same as the sum of both content and eyeball ISPs. We plot the value of $\phi_{\mathcal{T}}$ for $2 \leq |\mathcal{T}| \leq 8$. From Figure 4, we can make two observations. First, similar to the result of the CE model, the ratio of $\phi_C : \phi_{\mathcal{T}} : \phi_B$ is fixed when the ratio $|\mathcal{C}| : |\mathcal{T}| : |\mathcal{B}|$ is fixed. In other words, we have the scaling effect of the normalized Shapley value function as:

$$\varphi_B(b, t, c) = k\varphi_B(kb, kt, kc) \quad \forall k = 1, 2, 3, \dots$$

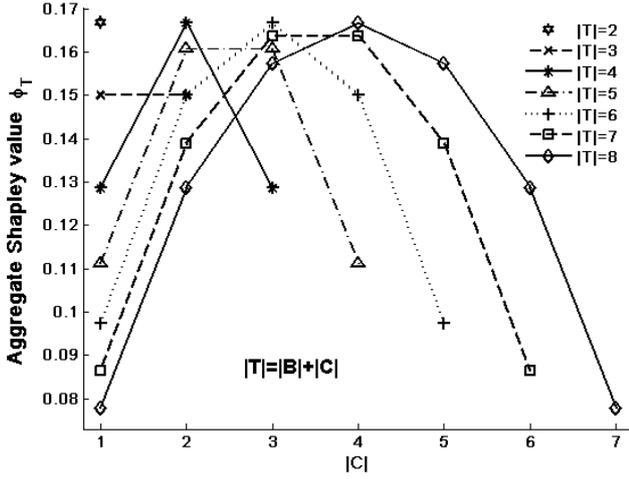


Figure 4: Aggregate Shapley value of transit ISPs.

For example, when we have only one ISP for each type, each ISP obtains one-third of the total revenue, i.e. $\varphi_B(1, 1, 1) = 1/3$. The aggregate Shapley for each group of ISP will keep the same, i.e. $\phi_B = v(\mathcal{N})/3$, as long as the sizes of the groups of ISPs increase proportionally. If they do not increase proportionally, we have the second observation. Even the number of transit ISPs keeps a constant as the sum other ISPs, its aggregate Shapley value changes as the sizes of content and eyeball ISPs vary. Each curve exhibits a reverted U-shape, where ϕ_T reaches its maximum when $|\mathcal{B}| = |\mathcal{C}|$. This result also coincides our intuition. When $|\mathcal{B}| = 1$ or $|\mathcal{C}| = 1$, the only eyeball or content ISP becomes crucial and shares a great amount of the total revenue. When the number of eyeball and content ISPs are evenly distributed, i.e. $|\mathcal{B}| = |\mathcal{C}|$, the impact of any of them leaving the system is minimized, which, at the same time, maximizes the value of transit ISPs.

4.3 Multiple Contents and Regions Model

In this section, we extend our previous result for multiple contents and multiple regions. The conservation of revenue in Equation (5) represents the aggregate revenue as summation of individual ISPs' revenue. Another way of decomposing the aggregate revenue is to separate different revenue sources from where the revenues are generated. We define $BP_r = \alpha_r X_r$ as the aggregate eyeball-side revenue generated in region r , and $CP_q^r = \beta_q X_r$ as the aggregate content-side revenue generated by providing content q in region r . As a result, we can decompose the total revenue in the network $v(\mathcal{N})$ as follows.

$$v(\mathcal{N}) = \sum_{r=1}^{|\mathcal{R}|} BP_r + \sum_{q=1}^{|\mathcal{Q}|} \sum_{r=1}^{|\mathcal{R}|} CP_q^r. \quad (6)$$

Intuitively, since any eyeball ISP in region r contributes to the eyeball-side revenue BP_r and any content ISP providing content q contributes to the content-side revenue CP_q^r for all region r , these ISPs should get a fair share of the specific

revenue they contributed to generate.

Theorem 3 (MULTIPLE CONTENTS/REGIONS MODEL). We consider a network with a set of content ISPs \mathcal{C} , a set of eyeball ISPs \mathcal{B} and a set of transit ISPs \mathcal{T} . Each content ISP C_i provides a set of contents $Q_i \subset \mathcal{Q}$ and each eyeball ISP B_j covers a set of regions $R_j \subset \mathcal{R}$. Both content and eyeball ISPs are connected to the transit ISPs by a complete bipartite graph. The Shapley value revenue for each ISP is in the following form:

$$\varphi_{B_j} = \sum_{r \in R_j} [\varphi_B(h_r, |\mathcal{T}|, |\mathcal{C}|)BP_r + \sum_{q \in \mathcal{Q}} \varphi_B(h_r, |\mathcal{T}|, h_q)CP_q^r],$$

$$\varphi_{T_k} = \sum_{r \in \mathcal{R}} [\varphi_T(h_r, |\mathcal{T}|, |\mathcal{C}|)BP_r + \sum_{q \in \mathcal{Q}} \varphi_T(h_r, |\mathcal{T}|, h_q)CP_q^r],$$

$$\varphi_{C_i} = \sum_{r \in \mathcal{R}} [\varphi_C(h_r, |\mathcal{T}|, |\mathcal{C}|)BP_r + \sum_{q \in Q_i} \varphi_C(h_r, |\mathcal{T}|, h_q)CP_q^r].$$

Theorem 3 shows that in a multiple contents/regions environment, the Shapley value revenue can be expressed as separable Shapley components of specific content-side and eyeball-side revenues. In particular, the eyeball-side revenue BP_r is not shared by the eyeball ISPs that are not covering region r and the content-side revenue CP_q^r is not shared by the content ISPs that are not providing content q . Each separated revenue is distributed among ISPs according to Theorem 2, using the normalized Shapley value functions φ_B, φ_T and φ_C . This result is a consequence of the linearity property of the Shapley value [12], which can be applied to more general topologies of the network.

4.4 Extension to General Internet Topology

In this subsection, we consider more general topologies of the network than the complete bipartite connections between ISPs before. Because transit ISPs model the Tier-1 ISPs in the Internet, we focus on the topologies where transit ISPs form a full mesh. However, most of the results apply to more general topologies.

To evaluate the Shapley value revenue for ISPs in a general topology, we first decompose the aggregate revenue according to Equation (6) in a similar manner as in Theorem 3. For each component of the aggregate revenue (i.e. a content-side revenue CP_q^r or an eyeball-side revenue BP_r), a subsystem can be derived to distribute it. For example, the system in Figure 3(d) can be decomposed to six subsystems in Figure 5. To construct the subsystems in Figure 5, we only include the eyeball ISPs in region r for BP_r and content ISPs that provide content q for CP_q^r . Because the topology is not a complete bipartite graph, some of the transit ISPs do not contribute for certain component of the total revenue. We need to eliminate these *dummy* transit ISPs as depicted in dotted circles in Figure 5.

Definition 3. An ISP i is dummy if $\Delta_i(v, \mathcal{S}) = 0$ for every $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$.

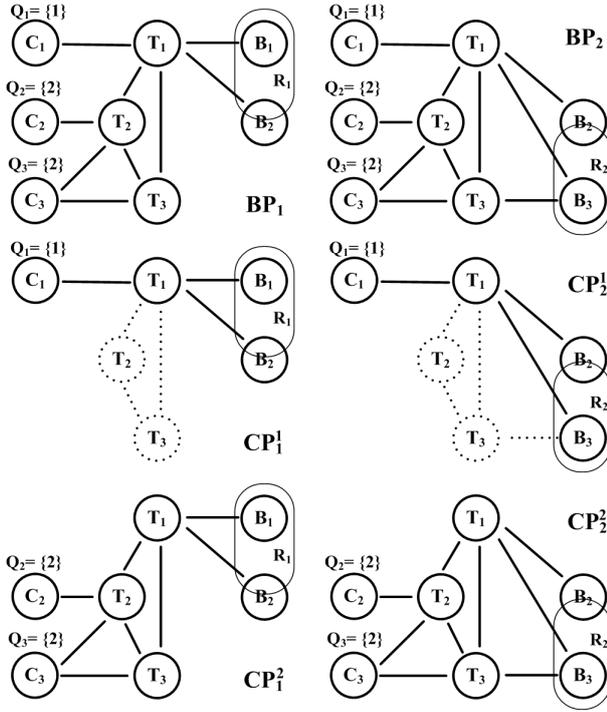


Figure 5: The decomposition of Shapley values.

The remaining problem is to solve the Shapley revenue distribution for each of the decomposed subsystems in Figure 5. Theorem 2 gives the closed-form solution for complete bipartite topologies. Although topological differences vary the Shapley revenues, we can evaluate the Shapley values via a dynamic programming procedure, if the system is *canonical*.

Definition 4. A system (\mathcal{N}, v) is *canonical*, if $v(\mathcal{S})$ is either 0 or $v(\mathcal{N})$ for every $\mathcal{S} \subseteq \mathcal{N}$.

In a canonical system, the total revenue $v(\mathcal{N})$ can be either wholly earned or lost by any coalition $\mathcal{S} \subseteq \mathcal{N}$. This implies that when any ISP leaves the system, the resulting topology does not segment the network as in Figure 2, otherwise partial revenue will be lost and $0 < v(\mathcal{S}) < v(\mathcal{N})$, which violates the canonical property. By having a full-mesh of the group of transit ISPs, each of the decomposed subsystems is indeed a *canonical* system. As a result, any coalition \mathcal{S} can obtain either the whole decomposed revenue or nothing.

Definition 5. An ISP i is called a *veto ISP*, if i belongs to all \mathcal{S} with $v(\mathcal{S}) > 0$.

Every veto ISP is essential for generating the revenue. If any veto ISP leaves the system, the worth of the remaining coalition becomes zero in a canonical system. For example, transit ISP T_1 is a veto ISP for the eyeball-side revenue BP_1 , because all eyeball ISPs have to go through it to obtain contents for customers. Now, we are ready to describe the dynamic programming procedure to solve the Shapley revenue for each of the decomposed canonical systems.

Theorem 4 (DYNAMIC PROGRAMMING EVALUATION). For any canonical system (\mathcal{N}, v) , we define $\{(\mathcal{S}, v) : \mathcal{S} \subset \mathcal{N}\}$

as the set of subsystems formed by any coalition \mathcal{S} of ISPs and $\varphi_i(\mathcal{S}, v)$ as the Shapley value of ISP i in the subsystem (\mathcal{S}, v) . The Shapley value revenue $\varphi_i(\mathcal{N}, v)$ for any ISP $i \in \mathcal{N}$ can be expressed as a function of the Shapley values from the subsystems $\{(\mathcal{S}, v) : \mathcal{S} \subset \mathcal{N}, |\mathcal{S}| = |\mathcal{N}| - 1\}$ as:

$$\varphi_i(\mathcal{N}, v) = \frac{1}{|\mathcal{N}|} \left[\sum_{j \neq i} \varphi_i(\mathcal{N} \setminus \{j\}, v) + v(\mathcal{N}) \mathbf{1}_{\{i \text{ is veto}\}} \right].$$

Theorem 4 shows that the Shapley values of a canonical system (\mathcal{N}, v) can be represented by the Shapley values of its subsystems (\mathcal{S}, v) that have one less cardinality of the number of ISPs. This result implies that we can build the Shapley values using a bottom-up dynamic programming approach that progressively calculates the Shapley values of the subsystems to form the Shapley values of the original canonical system. In practice, this procedure can also help calculate the Shapley value of a progressively developing system. For example, if all prior Shapley values are available, a new system with an ISP joining in can be calculated directly from the recursion equation in Theorem 4. Moreover, Theorem 2 can also be helpful in practice when a subsystem (\mathcal{S}, v) happens to have a complete bipartite topology.

4.5 Connectivity Effects on Shapley Revenues

Figure 4 compares the aggregate Shapley value for each group of ISPs, when the number of ISPs in each group changes. Here, we fix the number of ISPs in each group and explore how the aggregate values of each group of ISPs change when the interconnecting topology changes.

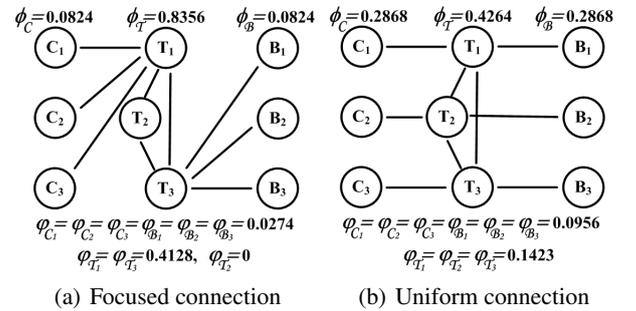


Figure 6: Two extreme ways to interconnect.

Figure 6 illustrates Shapley value distribution for ISPs when $|\mathcal{C}| = |\mathcal{T}| = |\mathcal{B}| = 3$ and each content or eyeball ISP only connects to one transit ISP. We can see that the Shapley values differ drastically, depending on how the content and eyeball ISPs are connected to the transit ISPs. Figure 6(a) shows the case when all content ISPs are connected to T_1 and all eyeball ISPs are connected to T_3 . Although T_2 becomes a dummy ISP, the group of transit ISPs possesses 83% of the total revenue. Figure 6(b) shows the case when all content and eyeball ISPs are connected to transit ISPs uniformly. In contrast, the group of transit ISPs only obtain 42%, half of the previous share, of the total revenue.

Figure 7 illustrates how the values of ϕ_B, ϕ_T and ϕ_C change reacting to the interconnecting link changes. Along the X-axis, we vary the degree of connectivity to the transit ISPs.

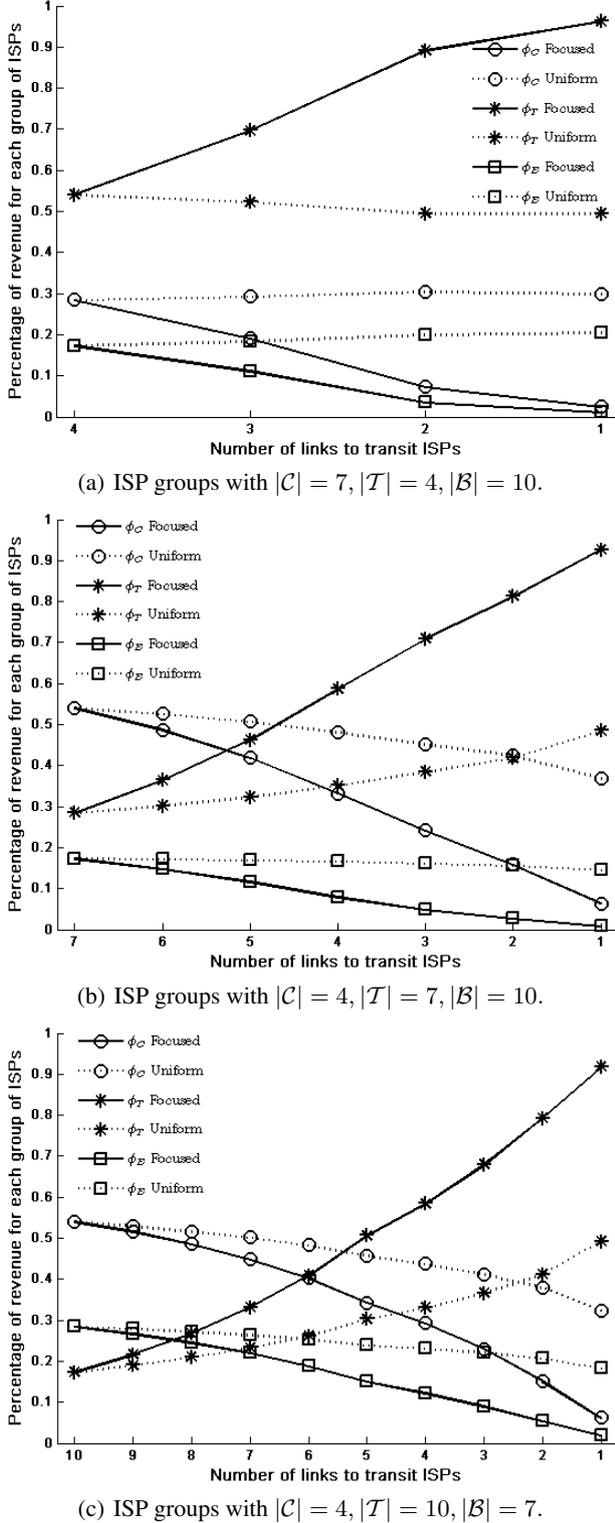


Figure 7: The Shapley value revenue for group of ISPs.

We start from the complete bipartite topology where each content and eyeball ISPs connect to all $|\mathcal{T}|$ transit ISPs. Then, we gradually decrease the number of transit ISPs each content and eyeball ISP connects to. We plot two types of connection topologies.

1. Focused (solid lines): For any degree of connectivity k , content ISPs connect on the first k transit ISPs and eyeball ISPs connect to the last k transit ISPs (as in Figure 6(a)).
2. Uniform (dotted lines): Content and eyeball ISPs connect to transit ISPs in a round-robin manner, where each transit ISP connects approximately the same number of content and eyeball ISPs (as in Figure 6(b)).

We observe that the value of the transit ISPs, ϕ_T , increases and both ϕ_C and ϕ_B decrease in general when the degree of connectivity decreases. Given any fixed degree of connectivity, ϕ_T is large when content and eyeball ISPs are focused on different transit ISPs, and is small when content and eyeball ISPs are connected uniformly. This general trend can be understood by considering the *effective* number of transit ISPs in the topology. Under the Focused topology, for example when $k = 1$, there are only two effective transit ISPs and other transit ISPs become dummy. Therefore, we can imagine the effective size of the transit ISPs as $|\mathcal{T}| = 2$, which results the large ϕ_T value for transit ISPs. In an extreme case in Figure 7(a) when $|\mathcal{T}|$ is small relative to $|\mathcal{C}|$ and $|\mathcal{B}|$, the value of ϕ_T might decrease a little bit when links are connected uniformly. Intuitively, under this scenario, the effective number of transit ISPs does not change much when the degree of connectivity decreases, because there are more content and eyeball ISPs than transit ISPs. Unlike the symmetric property shown on complete bipartite topologies, the degree of connectivity and the way how content and eyeball ISPs connect to transit ISPs strongly affect the value ϕ_T .

5. IMPLICATIONS

In the previous section, we developed the Shapley revenue distribution for ISPs under general Internet topologies. Although the Shapley value solution inherits multiple desirable property, the actual revenue distribution in the Internet might deviate from the Shapley value due to the inefficient bilateral agreements between ISPs. In this section, we discuss the implications derived from the Shapley value solution that may guide the establishment of bilateral agreements and the pricing structure for differentiated services. We start from a brief discussion of the stability of the Shapley value solution.

5.1 Stability

Stability concerns whether ISPs can form a coalition to earn more revenue than the aggregate Shapley revenue. If so, ISPs do not have incentives to cooperate all together and may deviate from the Shapley revenue distribution. Mathematically, the stability condition requires:

$$\sum_{i \in \mathcal{S}} \varphi_i(\mathcal{N}, v) \geq v(\mathcal{S}) \quad \forall \mathcal{S} \subseteq \mathcal{N}. \quad (7)$$

The above inequality requires the aggregate Shapley value of any coalition to be greater than or equal to the worth of the coalition; otherwise, the coalition will not cooperate with other ISPs under the Shapley value mechanism. All solutions that satisfy the above inequality are said to be in the *core* [19]. The Shapley value is known to be in the core [23] of a system if its worth function v is *convex*.

Definition 6. A worth function v is *convex* if for all coalition S and S' , $v(S) + v(S') \leq v(S \cup S') + v(S \cap S')$. Particularly, Lloyd Shapley proved that the marginal contributions $\Delta_i(v, S(\pi, i))$ defined in Equation (2) form the vertices of the core of a convex system [23]. The Shapley value, which is the average of vertices of the core, is located at the center of the core. Figure 8 illustrates the core (solid

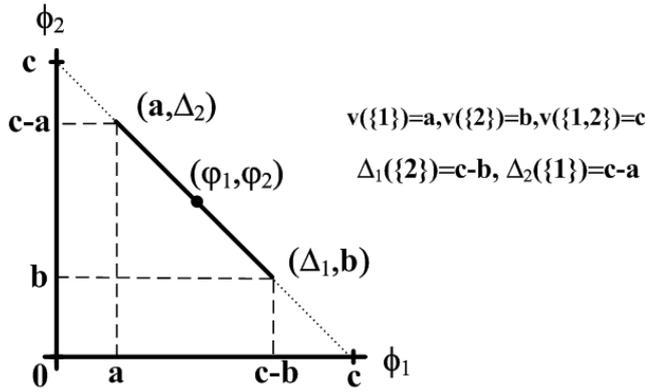


Figure 8: The core of a two-ISP example.

line segment) of a two-ISP example. The x-axis and y-axis represent the revenue distributed to ISP 1 and 2 respectively. The two vertices correspond to two marginal contributions: $\Delta_1(\{2\})$ and $\Delta_2(\{1\})$. The Shapley solution is located in the middle of the core. Notice that if ISP 1 gets less than a , it will not cooperate; if it gets more than $c - b$, ISP 2 will not cooperate, because ISP 2's gain is less than b . The convexity condition is pretty loose, because in reality, cooperation normally achieves higher values to the bigger coalition than the sum of individual values, i.e. $v(S) + v(S') \leq v(S \cup S')$. Notice that some systems naturally have an empty core, which means there does not exist any solution for Inequality (7).

5.2 Implications for Bilateral Agreements

The Shapley value solution suggests a value chain illustrated in Figure 9. End-payments flow into the network either from the content-side or the eyeball-side. Each group of ISPs retains a proportion, i.e. the Shapley revenue of the group, of the revenue and forwards the remaining along the network. In practice, ISPs negotiate bilateral settlements. Huston [9] concluded that the zero-dollar peering and the customer/provider relationships were the only stable models for the Internet at the '90s. The effective revenue distribution resulted from these bilateral agreements probably deviate from that of the Shapley value distribution. We ask why these bilateral agreements were stable and how close they

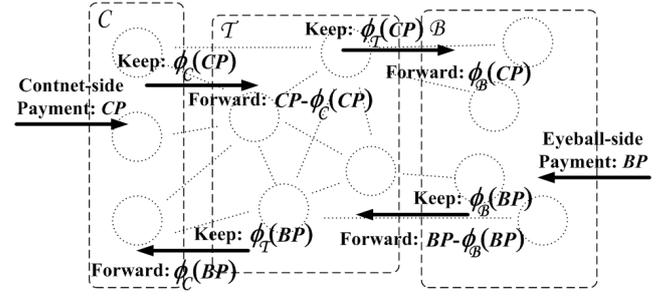


Figure 9: The value chain to implement Shapley revenue.

were to the Shapley solution. Figure 10 illustrates the sce-

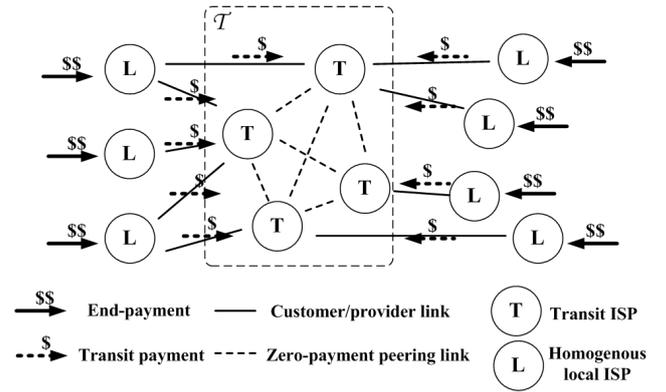


Figure 10: Traditional ISP structure with homogeneous local ISPs and transit ISPs.

nario when local ISPs were still homogeneous and the end-to-end traffic patterns exhibited symmetry at the '90s. We can imagine that local ISPs were not specialized to be content or eyeball ISPs. They obtain end-payment both from content-providers and end-customers. In the Shapley value framework, ISPs need to exchange different payments. Each local ISP will forward more money to transit ISPs than they receive from transit ISPs. Effectively, the net money exchange would be from local ISPs to transit ISPs. Due to the symmetric traffic pattern, the net money exchange between transit ISPs would be close to zero. This result coincides with zero-dollar peering and the customer/provider relationships established from bilateral agreements. Although the exact revenue distribution might deviate from the Shapley value solution, we conjecture that the resulting revenue distribution was close enough to the Shapley value solution so that it was in the core, and thus stable.

Faratin et al. [5] observed that due to the erosion of homogeneity of ISPs, *specialized* ISPs (content and eyeball) have emerged as well as a new type of bilateral agreement: *paid-peering*. Paid-peering is identical to zero-dollar peering in terms of traffic forwarding, except that one party needs to pay another. By applying the Shapley revenue distribution to the Content-Transit-Eyeball model, we find the justification of the existence of paid-peering between transit ISPs. Figure

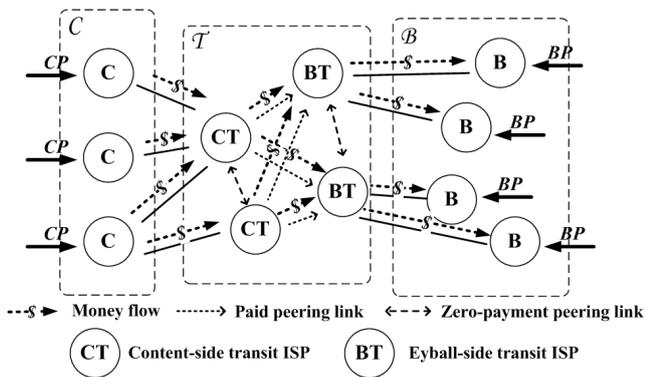


Figure 11: Shapley value implied money exchange.

11 illustrates a scenario where the content ISPs connect to a set of Content-side Transit (CT) ISPs and eyeball ISPs connect to a set of Eyeball-side Transit (BT) ISPs. Assume that the eyeball-side payments are relatively small compared to the content-side payments, because BP s are based on fixed monthly payments from end-customers and CP s are growing with the Internet-related businesses. After netting the exchange of payments along the value chain in Figure 9, we show the net bilateral money flows that result the Shapley value solution in Figure 11. We observe that content ISPs obtain content-side payments and pay CT ISPs. This is the same customer/provider relationship as before. However, the zero-dollar peering relationship does not happen between all pairs of transit ISPs. Notice that CT ISPs need to forward the content-side value towards the eyeball-side, which turn out to be the paid-peering relationship emerged with heterogeneous ISPs. One unconventional observation is that, eyeball ISPs need to receive compensations from content-side through BT ISPs. This implies that transit ISPs should pay eyeball ISPs, which creates a reverse customer/provider relationship. In reality, this reverse customer/provider settlement rarely happens, because transit ISPs do not pay their customer ISPs. From these implied bilateral relationships, we realize that the current practice of bilateral agreements may probably reach a solution that deviates from the theoretic Shapley solution severely. Consequently, this revenue solution locates outside the core, and thus is unstable. We conjecture that Level 3's de-peering with Cogent might be the result of failing to implement an appropriate paid-peering agreement as implied by the Shapley solution. Also, the failure to maintain the reverse customer/provider relationship might be the reason that the incumbent last-mile service-providers complain that they are not paid enough to recover their investments. Consequently, they want to create service differentiations to recoup extra profit. This naturally leads to the debate of *network neutrality*.

5.3 Implications for Differentiated Services

The centerpiece of the network neutrality debate is the necessity to impose potential regulatory enforcements, by which telephony companies have been regulated, on the In-

ternet. The proponents [3, 25] criticized the discriminatory behavior by ISPs, believing that it harms the productivity, innovation and end-to-end connectivity of the Internet. However, the opponents [11] advocated that offering premium service stimulates innovation on the edges of the network. Musacchio et al. [18] show that different parameters, e.g. advertising rate and end user price sensitivity, influence whether a neutral or non-neutral regime achieves a higher social welfare.

As we discussed before, bilateral agreements that severely deviate from the Shapley revenue distribution will cause unstable interconnections among ISPs. Similarly, even though differentiated services can be shown to be beneficial to the network and end users, without an appropriate revenue distribution mechanism, ISPs do not have the incentive to architect the cooperative provisioning of such services. As a generic revenue sharing mechanism, the Shapley value solution can also be used for encouraging ISPs to participate and fairly share extra revenue amongst participating ISPs. Here, we illustrate two potential differentiated services and the implied compensation structure for the supporting ISPs.

5.3.1 Supporting gaming services

The booming online gaming industry has brought huge revenues. The current 4 billion worth of the global online game market is expected to triple in the next five years according to Strategy Analytics's outlook for the market. ABI Research predicts that the online game segment of the game industry will grow by 95% each year until 2011, when it becomes the dominating force in the market. In order to support networked games with required low latency and accurate synchronization, network providers need to provide differentiated services for game-providers. However, an appropriate compensation structure is crucial for providing incentives for network providers to establish such levels of quality service to support game applications.

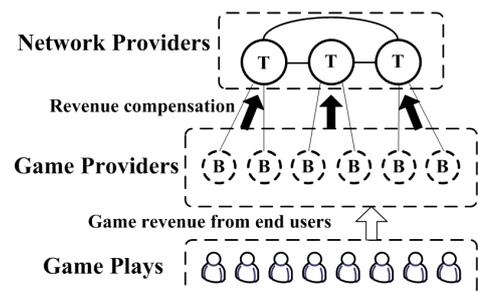


Figure 12: Compensation structure for game services.

Figure 12 shows the compensation structure implied by the Shapley value solution for game services. By providing gaming applications to players, game providers (dotted circles) obtain extra revenue and can be considered as eyeball ISPs who serve end-users. Network providers (solid circles) can be considered as transit ISPs, providing interconnections between the customers from different game-providers. Due to the symmetry of the network traffic pattern for network

games, the compensation structure is similar to the “customer/provider and zero-dollar peering” structure in Figure 10: game providers need to compensate network providers for supporting the new service and network providers connect to one another with zero-dollar peering agreements.

5.3.2 Supporting financial services

Another potential application is financial services across the Internet. Many financial applications require a low latency to retrieve accurate realtime data, e.g. realtime stock or option quotes; others require guaranteed network services, e.g. trading transactions. High (or asymmetric) latencies can make these applications extremely vulnerable, as a few milliseconds here or there can translate to billions of dollars with automated trading. Network security is another major concern for providing secured financial transactions on the network. In order to implement more robust and secure protocols across the Internet, cooperations among ISPs might be needed to support low latency and to prevent malicious attackers and information stealers.

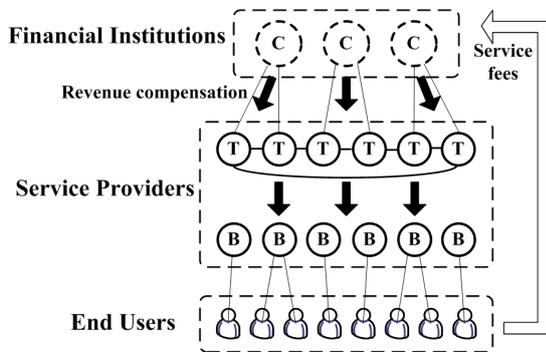


Figure 13: Implied compensation structure for financial services.

Figure 13 shows the compensation structure implied by the Shapley value solution for secure financial services. On the top, the financial institutions (dotted circles) earn revenue from customers by providing the online financial services. They can be considered as the content providers who obtain revenue by attracting business from customers’ ability to access to the Internet. Service providers (solid circles) cooperatively implement secure protocols to maintain the interaction between customers and financial institutions. ISPs directly connect to customers are eyeball ISPs and intermediate ISPs are transit ISPs. Due to the asymmetric traffic characteristic, the compensation structure is similar to the one shown in Figure 11. Financial institutions compensate transit ISPs as in a customer/provider relationship. Transit ISPs need to compensate eyeball ISPs for the reverse customer/provider relationship. Paid-peering relationship might also exist if multiple levels of transit ISPs are in the network.

6. RELATED WORK

Our previous work [12] proposed a clean-slate revenue distribution mechanism for ISPs, based on the Shapley value

in a general network setting. We showed that under the Shapley value mechanism, selfish ISPs have incentives to perform globally optimal routing and interconnecting decisions to reach an equilibrium that maximizes both individual profit and global social welfare. Due to the multi-lateral nature of the mechanism and the exponential complexity of the Shapley value, how to implement and use the Shapley value solution was an unsolved problem. Our first attempt [13] to model a detailed Internet structure was limited to the Content-Eyeball model introduced by Faratin et al. [5]. In this paper, we extend our model to include a third class of ISPs: the transit ISPs. We generalize all results in [13] as special cases of a multiple contents/regions model (Theorem 3). We explore the closed-form Shapley solution under structured topologies and develop a dynamic programming procedure to compute the Shapley solution for general topologies.

Bailey [1] and Huston [9] started exploring the interconnection settlements of the ISP in the ’90s. Huston [9] and Frieden compared the existing Internet settlement models with that of the telecommunication industry’s. Due to the irregularity of the Internet structure, none of the traditional telecommunication settlement model can be brought into the Internet. Based on empirical evidences, Huston conjectured that the zero-dollar peering and the customer/provider relationships were the only stable models for the Internet at the time. Faratin et al.’s recent work on ISP settlement [5] exhibits interconnection disputes in the Internet and observes the emergence of *paid-peering* relationship between ISPs. Our work explores the bilateral relationship implied by the Shapley value solution. Our result validates that under the symmetric traffic pattern and the homogeneity of the ISPs, zero-dollar peering and the customer/provider relationships can create a stable equilibrium that is close to the Shapley value. Under the CTE model, the Shapley value solution also validates existence of paid-peering relationship between transit ISPs. Moreover, it also suggests that a reversed customer/provider relationship should exist between transit and eyeball ISPs. Our result explains the origin of failures of current bilateral agreements, e.g. de-peering and the emergence of network neutrality debate.

Gao [8] proposed a relationship-based model for ISPs and categorized the interconnection relationship by provider-to-customer, peer-to-peer and sibling-to-sibling links. However, Battista et al. [2] experimented on AS relationships and observed violations of the valley-free property [8] from BGP routing tables. Our work treats ISPs as cooperative entities that form coalitions to share revenue. The reverse customer/provider relationship implied from the Shapley value solution under the CTE model can explain the violations of valley-free property found in the AS-paths.

The network neutrality debate [25, 3, 7] started when discriminatory practices, e.g. selectively dropping packets, were found with broadband provider and cable operators. Crowcroft [3] reviewed technical aspects of network neutrality and concluded that network neutrality should not be engineered. Both sides of the debate are concerned about whether differenti-

ated services should be provided in the Internet. Musacchio et al. [18] derived different regions that network neutrality can be good or bad to the whole network. Our work provides an orthogonal thought about the differentiated services: the appropriateness of providing differentiated services depend on a suitable pricing structure for the ISPs that provide the service. We propose that the Shapley solution can be used as the pricing structure to encourage individual incentives and increase social welfare.

Originated from microeconomics theory [17], game theory [19] has been used to address pricing [22] and incentive problems [15] in networking areas. Unlike the majority of noncooperative game models, the Shapley value [20] originates from *coalition games* [19] that model the cooperative nature of groups. Eyal Winter's survey [24] provides a through investigation on the Shapley value and its properties.

7. CONCLUSION

In this paper, we explore the Shapley value solution for a detailed Internet model with three classes of ISPs: content, transit and eyeball. We derive closed-form solutions for structured topologies and a dynamic programming procedure to evaluation solutions under general topologies. In particular, we prove that a complex system with multiple revenue sources from different contents and regions can be decomposed by their inelastic components of content-side and eyeball-side revenues. Because the Shapley value often locates at the center of the core, which contains all stable revenue distribution solutions, we use the Shapley value solution as a benchmark to validate the stability of bilateral agreements used in the past and current Internet. We find that because the symmetry of traffic flow and the homogeneity of ISPs, traditional zero-dollar peering and customer/provider relationship can create stable solutions that are close to the Shapley value solution. However, when ISPs exhibit heterogeneity and traffic flows are mainly from content-side to eyeball-side, the solutions from bilateral agreements severely deviate from the Shapley value solution, which exhibits a paid-peering relationship between transit ISPs and a reverse customer/provider relationship between transit and eyeball ISPs. We conjecture that many of the failures of bilateral agreements are due to the lack of implementing these paid-peering and reverse customer/provider relationship via bilateral agreement. Finally, we propose to use the Shapley value solution as the pricing structure for differentiated services so that ISPs will be encouraged to fairly share newly brought revenues and improve the Internet services. We believe that these results can be useful for settling bilateral disputes and for regulatory institutions to regulate the industry.

8. REFERENCES

- [1] J. P. Bailey. The economics of Internet interconnection agreements. *Internet economics*, pages 155 – 168, 1997.
- [2] G. D. Battista, T. Erlebach, A. Hall, M. Patrignani, M. Pizzonia, and T. Schank. Computing the types of the relationships between autonomous systems. *IEEE/ACM Transactions on Networking*, 15(2):267–280, April 2007.
- [3] J. Crowcroft. Net neutrality: the technical side of the debate: a white paper. *ACM SIGCOMM Computer Communication Review*, 37(1), January 2007.
- [4] G. Demange and M. Wooders. *Group formation in economics: networks, clubs, and coalitions*. Cambridge University Press, Cambridge, 2005.
- [5] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr. Complexity of Internet interconnections: Technology, incentives and implications for policy. *The 35th Research Conference on Communication, Information and Internet Policy (TPRC)*, 2007.
- [6] R. Frieden. Without public peer: the potential regulatory and universal service consequences of Internet Balkanization. *Virginia Journal of Law and Technology*, 3, 1998.
- [7] R. Frieden. Network neutrality or bias? – handicapping the odds for a tiered and branded Internet, 2006.
- [8] L. Gao. On inferring autonomous system relationships in the internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, December 2001.
- [9] G. Huston. *ISP Survival Guide: Strategies for Running a Competitive ISP*. John Wiley and Son, New York, 1999.
- [10] M. O. Jackson. Allocation rules for network games. *Game Theory and Information from EconWPA*, 2003.
- [11] M. Jamison and J. Hauge. Getting what you pay for: Analyzing the net neutrality debate. *University of Florida, Department of Economics, PURC Working Paper*, 2007.
- [12] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. Internet Economics: The use of Shapley value for ISP settlement. *Proceedings of 2007 ACM Conference on Emerging network experiment and technology (CoNEXT 2007)*, December 2007.
- [13] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. Interconnecting Eyeballs to Content: A Shapley Value Perspective on ISP Peering and Settlement. *Proceedings of 2008 ACM Network Economics (NetEcon)*, August 2008.
- [14] R. T. B. Ma, D. M. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. The Shapley revenue for content, transit and eyeball ISPs. *technical report, Columbia University, EE Department, July 2008*. [Online] Available: <http://dn-pubs.cs.columbia.edu/citation/paperfile/169/tech08.pdf>.
- [15] R. T. B. Ma, S. C. Lee, J. C. Lui, and D. K. Yau. Incentive and service differentiation in P2P networks: A game theoretic approach. *IEEE/ACM Transactions on Networking*, 14(5):978–991, October 2006.
- [16] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, kc claffy, and A. Vahdat. Lessons from three views of the Internet topology. *CAIDA report*, 2005.
- [17] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- [18] J. Musacchio, G. Schwartz, and J. Walrand. Network neutrality and provider investment incentives. *Asilomar Conference*, pages 1437–1444, November 2007.
- [19] M. J. Osborne and A. Rubinstein. *A course in game theory*. The MIT Press Cambridge, Massachusetts, 1994.
- [20] A. Roth. *The Shapley value: Essays in honor of Lloyd S. Shapley*. Cambridge University Press, Cambridge, 1988.
- [21] T. Roughgarden and E. Tardos. How bad is selfish routing? In *IEEE Symposium on Foundations of Computer Science*, pages 93–102, 2000.
- [22] S. Shakkottai and R. Srikant. Economics of network pricing with multiple ISPs. In *Proceedings of IEEE Infocom*, 2005.
- [23] L. Shapley. Cores of convex games. *International Journal of Game Theory*, 1:11–26, 1971.
- [24] E. Winter. *The Shapley Value, in The Handbook of Game Theory*. R. J. Aumann and S. Hart, North-Holland, 2002.
- [25] T. Wu. Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law*, 141, 2005.