

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

Richard T. B. Ma, Dah Ming Chiu, *Fellow, IEEE*, John C. S. Lui, *Fellow, IEEE, ACM*, Vishal Misra, *Member, IEEE*, and Dan Rubenstein, *Member, IEEE*

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 profit 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 values 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.

Index Terms—Coalition games, Internet service provider (ISP) settlement, network economics, network neutrality, Shapley value.

I. INTRODUCTION

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 [25] and interconnecting [8], 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 being 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 ongoing 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 [2], [11] 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 [11]. Today, because of the heterogeneity in ISPs, simple peering agreements are not always satisfactory to all parties involved, and *paid peering* [7] 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* [4], [9], [30] debate, which argues the appropriateness of providing service/price differentiations in the Internet.

Our previous work [17] explored the application of *Shapley value* [24], [29], a well-known economic concept originated from coalition games [5], [12], [23], 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. These results demonstrate the viability of a Shapley value mechanism under the ISP profit-sharing context.

In this paper, we explore the *Shapley value profit 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.* [7] view today’s Internet as containing two classes of ISPs: *eyeball* and *content*. Eyeball ISPs, such as Time Warner Cable, Comcast, and Verizon ADSL, specialize in delivery to hundreds of thousands of residential users, i.e., supporting the last-mile connectivity. Content ISPs specialize in providing hosting and network access for end-users and commercial companies that offer contents, such as Google, Yahoo!, and YouTube. Typical examples are Cogent and content distribution networks (CDNs) like Akamai. Our previous work [15], [16] explored the Shapley value revenue distribution based on this content–eyeball (CE) model. This paper starts with the CE model and extends it to

Manuscript received January 20, 2010; revised July 14, 2010 and October 08, 2010; accepted October 13, 2010; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor C. Dovrolis. Date of publication November 09, 2010; date of current version June 15, 2011.

R. T. B. Ma is with the Department of Computer Science, National University of Singapore, Singapore 117418, Singapore, and also with the Advanced Digital Sciences Center, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: tbma@comp.nus.edu.sg).

D. M. Chiu is with the Department of Information Engineering, Chinese University of Hong Kong, Shatin, NT, Hong Kong (e-mail: dmchiu@ie.cuhk.edu.hk).

J. C. S. Lui is with the Department of Computer Science and Engineering, Chinese University of Hong Kong, Shatin, NT, Hong Kong (e-mail: cslui@cse.cuhk.edu.hk).

V. Misra is with the Department of Computer Science, Columbia University, New York, NY 10027 USA (e-mail: misra@cs.columbia.edu).

D. Rubenstein is with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: danr@cs.columbia.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2010.2089533

consider profit distribution and to include a third class: transit ISPs. Transit ISPs model the Tier-1 ISPs, such as Level 3, Qwest, and Global Crossing, which provide transit services for other ISPs and naturally form a full-mesh topology to provide the universal accessibility of the Internet. The three types of ISPs are more of a conceptual classification rather than a strict definition. For example, Akamai might not be regarded as an ISP in a strict sense. However, Google might be considered as a content ISP because it deployed significant network infrastructure and contributed more than 5% of the total Internet traffic in 2009 as reported in [14]. Also, an ISP might play multiple roles, e.g., Verizon can be both an eyeball and a transit ISP after acquiring Tier-1 UUNET (AS 701). Many Tier-1 ISPs are also providing CDN services to service content providers. Other classification models can also be found in literature, e.g., Dhamdhare *et al.* [6] classifies the ISPs as five types. However, our results are not restricted to our content–transit–eyeball (CTE) model. Our new results are the following.

- We obtain closed-form Shapley revenue and cost solutions for ISPs in the CE and the CTE models under bipartite topologies (Theorems 1, 2, 5, and 6).
- We generalize the closed-form Shapley revenue for multiple contents/regions environments where *inelastic* components can be decomposed linearly (Theorem 3).
- 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: 1) the aggregate revenue can be decomposed by content-side and eyeball-side components; and 2) the costs can be decomposed with respect to individual ISPs. Each revenue/cost component can be distributed as a Shapley value of a *canonical* subsystem to each ISP that contributed in the coalition.
- Through the Shapley value solution, we justify: 1) why the zero-dollar peering and customer/provider bilateral agreements could be stable in the early stage of the Internet; 2) why, besides operational reasons, paid-peering has emerged; 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 proved 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 pricing regulations, and for developers to negotiate and provide differentiated services on top of the current Internet.

II. SHAPLEY VALUE AND PROPERTIES

We follow the notation in [17] and briefly introduce the concept of Shapley value and its use under our ISP profit distribution context. We consider a network system comprising 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 subnetwork that might be able to provide partial services to their users. We denote v as the *worth function*, which measures the monetary benefits produced by the subnetworks formed by all coalitions. In particular, for any coalition \mathcal{S} , $v(\mathcal{S})$ denotes the profit (revenue minus cost) generated by the subnetwork formed by the set of ISPs \mathcal{S} , defined as

$$v(\mathcal{S}) = \hat{v}(\mathcal{S}) - \bar{v}(\mathcal{S}) \quad (1)$$

where $\hat{v}(\mathcal{S})$ and $\bar{v}(\mathcal{S})$ are the revenue and cost components of the worth function. 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], [29], the Shapley value serves as an appropriate mechanism for ISPs to share profit.

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 above definition 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 desirable efficiency and fairness properties [17].

As showed in [17], the Shapley value mechanism also induces global Nash equilibria that are globally optimal for routing and interconnecting. However, calculating the Shapley value often involves exponential time complexity [1]. By the additivity property [27] of the Shapley value, we can linearly decompose the Shapley value profit into a Shapley revenue component $\hat{\varphi}_i$ and a Shapley cost component $\bar{\varphi}_i$

$$\varphi_i(\mathcal{N}, v) = \varphi_i(\mathcal{N}, \hat{v}) - \varphi_i(\mathcal{N}, \bar{v}) = \hat{\varphi}_i(\mathcal{N}) - \bar{\varphi}_i(\mathcal{N}).$$

In this paper, we focus on the calculation of the Shapley revenue (Section IV), the Shapley cost (Section V), and the implications (Section VI) derived from the Shapley solution.

III. NETWORK MODEL

Faratin *et al.* [7] 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 all 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 $Q_i \subseteq \mathcal{Q}$ of the contents. 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 where the set of eyeball ISP \mathcal{B} provides Internet services to residential users. Each eyeball ISP B_j serves a subset $R_j \subseteq \mathcal{R}$ of the regions. We assume that each region $r \in \mathcal{R}$ has a fixed user population of size X_r . Each user

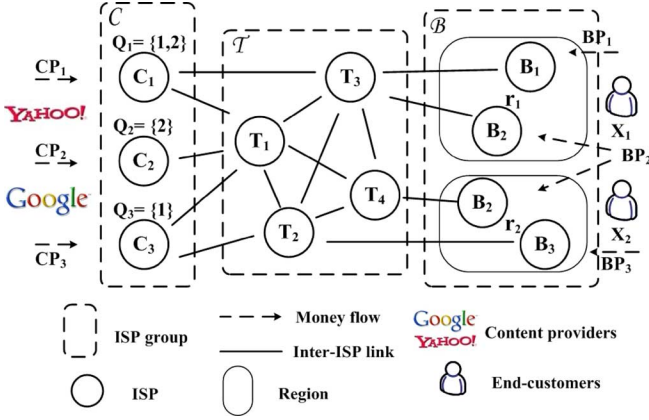


Fig. 1. Content–transit–eyeball ISP model.

chooses one of the eyeball ISPs serving the region for Internet service; therefore, each eyeball ISP B_j attracts and serves a portion x_j^r (equals zero if B_j does not serve 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 revenue from content providers to ISP C_i and BP_j as the eyeball-side revenue from residential users to ISP B_j .

Fig. 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 $Q_3 = \{1\}$. The regions that the three eyeball ISPs serve are $R_1 = \{1\}$, $R_2 = \{1, 2\}$, and $R_3 = \{2\}$.

Our network model represents the real Internet ISP structure. The justification comes from the study by the Cooperative Association for Internet Data Analysis (CAIDA) [20]. 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 AS path patterns are captured in our content–transit–eyeball model: Data sources originate from content ISPs, go through either one or two transit ISPs, and reach eyeball ISPs. The full-mesh topology of the transit ISPs is also true because the Tier-1 ISPs provide the universal accessibility of the Internet in reality [7].

A. Revenue and Cost Model

We define each eyeball-side revenue BP_j as

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

where α_r is the monthly charge in region r . The eyeball-side revenue is the aggregate service charge received from the residential users 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. Similarly, we define each content-side revenue CP_i as

$$CP_i = \sum_{q \in Q_i} \beta_q \sum_{r \in \mathcal{R}} \gamma_{ir}^q X_r \quad \forall C_i \in \mathcal{C} \quad (4)$$

where β_q is the average per-content revenue generated by delivering content q , and γ_{ir}^q is the fraction of users in region r that download content q from the content ISP C_i . Content ISPs receive payments for the data traffic they carry as well as the data services they provide for the customers of the content providers. The parameter β_q is affected by three factors. First, β_q might depend on the size of the content, which further implies the traffic volume and the corresponding transit cost it induces for the content ISPs. Second, β_q might depend on the specific service class for serving content q , which results in different quality of services and availabilities for the content. Third, β_q might also reflect the price differentiation imposed/provided by the content ISPs. For example, nonprofit organizations, e.g., universities, might be charged less than commercial companies, and large content providers might be able to negotiate lower per-bit price due to the economics of scale of their traffic. Detailed models for differentiated, service-based, or volume-based pricing are out of the scope of this paper. We denote γ_r^q as the fraction of users that download content q in region r defined as $\gamma_r^q = \sum_{i \in \mathcal{C}} \gamma_{ir}^q$. In practice, γ_r^q reflects the popularity of content q in region r . When the content is popular, a larger fraction of the population X_r will request for it, generating more revenue for the content ISPs.

We denote BL_j , TL_k , and CL_i as the cost of eyeball ISP j , transit ISP k , and content ISP i , respectively. In practice, the cost of an ISP includes the investments to build infrastructures and operation expenses. Rather than restricting the cost in a specific form, we consider the average cost of an ISP over a certain period of time. The cost can be calculated as a summation of evenly amortized investments plus the average operation expenses. The operation cost depends on the traffic volume an ISP carries. Therefore, it might also affect the parameter β_q for the content ISPs. For variable routing costs, please refer to a detailed model in our previous work [17].

B. User Demand Assumptions

We model how residential users choose to attach to different eyeball ISPs under various coalitions.

Definition 3: The demand of a residential user y for ISP B_j is *elastic* to ISP $B_{j'}$ if y would use $B_{j'}$ as an alternative when B_j becomes unavailable.

We assume that the intraregion user demand is *elastic* and the interregion user demand is *inelastic*. The elastic user demand within a region models a competitive market of eyeball ISPs that provide substitutive services to users. Therefore, users would not be sticky to a certain eyeball ISP within a region. On the other hand, the inelastic user demand across regions models the physical limitation of residential users to choose eyeball ISPs in a different geographical territory. By the above assumptions, the aggregate user demand in each region r is the fixed population X_r , regardless of the number of eyeball ISPs serving the region (as long as the number is greater than zero). Consequently, when a new eyeball ISP comes into a region, some of the users of the original ISPs will shift to the new ISP; when an existing eyeball ISP leaves a region, its users will shift to the remaining ISPs.

Conceptually, we can imagine that the demand of a particular content q is “elastic” to the set of content ISPs that provide q because users could be directed to any of these content ISPs for




Fig. 2. Segmented network results revenue loss.

downloading. Similarly, the demands of two different contents can be thought of as “inelastic” between the two sets of content ISPs because users have to download from the set of content ISPs that provides the particular content.

C. Conservation of Revenue

In general, the aggregate revenue generated by the entire network is a constant

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

However, the network might have revenue loss if the network is segmented so that some users cannot reach all the contents provided by all the content ISPs.

Fig. 2 illustrates an example where users, attached to one of the eyeball ISPs, can only access one of the content ISPs. The content-side revenue will be less than $\sum_{q \in \mathcal{Q}, r \in \mathcal{R}} \beta_q \gamma_r^q X_r$. However, if the two transit ISPs are connected, the total revenue will follow (5). In this paper, we assume that (5) holds under any topological change in the network (e.g., interconnection changes or ISP arrival/departure). The justification of this assumption is that, in practice, the set of transit ISPs is made up of the Tier-1 ISPs that always form a full mesh [7], [20]. Therefore, $\hat{v}(\mathcal{N})$ in (5) really defines the total revenue generated by the Internet.

Remark: The market price for accessing the Internet α_r might depend on the number of eyeball ISPs in the region r . This implies that with various grand coalition \mathcal{N} , $\hat{v}(\mathcal{N})$ can be different. The dynamics of market prices and the entry and exit of ISPs can be an orthogonal aspect of the model. However, we will derive our Shapley solutions in terms of a percentage of the value of the grand coalition $\hat{v}(\mathcal{N})$.

D. Coalitional Cost

Based on the demand assumptions, any coalition $\mathcal{S} \subset \mathcal{N}$ will induce revenue of $\hat{v}(\mathcal{S})$. We denote $\check{\mathcal{S}} \subseteq \mathcal{S}$ as the subset of nondummy ISPs in coalition \mathcal{S} , defined as

$$\check{\mathcal{S}} = \{s \in \mathcal{S} : s \text{ is not dummy in } (\mathcal{S}, \hat{v})\}.$$

The cost $\bar{v}(\mathcal{S})$ of the coalition \mathcal{S} is defined as the following:

$$\bar{v}(\mathcal{S}) = \sum_{s \in \check{\mathcal{S}}} L_s = \sum_{B_j \in \check{\mathcal{S}}} BL_j + \sum_{T_k \in \check{\mathcal{S}}} TL_k + \sum_{C_i \in \check{\mathcal{S}}} CL_i. \quad (6)$$




Fig. 3. Progressively developed models. (a) Content-eyeball model. (b) Content-transit-eyeball model. (c) Multiple contents/regions. (d) General topology.

The cost of a coalition focuses on nondummy ISPs. This also naturally includes the case where a coalition of disjoint ISPs cannot generate revenue and profit.

IV. 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, and then extend it with transit ISPs. After that, we further extend the model for multiple contents and regions. Finally, we explore more general Internet topologies under the previous models. Fig. 3 illustrates the four models we are going to discuss.

A. Content-Eyeball (CE) Model

The CE model follows Faratin *et al.* [7]. We focus on the single-content/single-region model for the time being, assuming that every content ISP provides the same content q and every eyeball ISP serves the same region r . As illustrated in Fig. 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 $\hat{v}(\mathcal{N}) = (\alpha_r + \beta_q \gamma_r^q) X_r$. We define $\hat{\phi}_{B_j}$ and $\hat{\phi}_{C_i}$ as the Shapley value revenue distributed to B_j and C_i , respectively, and $\phi_B = \sum_{B_j \in \mathcal{B}} \hat{\phi}_{B_j}$ and $\phi_C = \sum_{C_i \in \mathcal{C}} \hat{\phi}_{C_i}$ 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 [18].

Theorem 1 (The Shapley Revenue for the CE Model): We consider a set \mathcal{C} of content ISPs, providing one content and a set \mathcal{B} of eyeball ISPs, serving one region. Under the CE model with

a complete bipartite graph topology, the Shapley value revenue of each ISP is the following:

$$\begin{aligned}\hat{\phi}_{B_j}(|\mathcal{B}|, |\mathcal{C}|) &= \frac{|\mathcal{C}|}{|\mathcal{B}||\mathcal{N}|} \hat{v}(\mathcal{N}) \quad \forall B_j \in \mathcal{B} \\ \hat{\phi}_{C_i}(|\mathcal{B}|, |\mathcal{C}|) &= \frac{|\mathcal{B}|}{|\mathcal{C}||\mathcal{N}|} \hat{v}(\mathcal{N}) \quad \forall C_i \in \mathcal{C}.\end{aligned}$$

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 different type. In particular, the aggregate Shapley value revenue of both types of ISPs is inverse-proportional to the number of ISPs of each type, i.e., $\phi_{\mathcal{C}} : \phi_{\mathcal{B}} = |\mathcal{B}| : |\mathcal{C}|$.

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

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

This marginal aggregate revenue satisfies

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

Corollary 1 measures the marginal revenue loss of the set of eyeball ISPs for losing one of the content ISPs. Because the Shapley revenue in Theorem 1 is symmetric between content and eyeball ISPs, a similar marginal revenue result can be derived by considering any de-peering of an eyeball ISP. When $|\mathcal{C}| = 1$, Corollary 1 tells that the marginal revenue for the group of eyeball ISPs is $-\phi_{\mathcal{B}}$, 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 $\Delta_{\phi_{\mathcal{B}}} = -\frac{1}{|\mathcal{C}|^2}\phi_{\mathcal{B}}$, which means that by disconnecting an additional content ISP, this eyeball ISP is going to lose $1/|\mathcal{C}|^2$ of its revenue. The CE model gives a good sense of the Shapley value revenue distribution before it gets more complicated. Nevertheless, we will see that more detailed models show similar revenue-sharing features as the basic model.

B. Content–Transit–Eyeball (CTE) Model

The CTE model, as illustrated in Fig. 3(b), extends the CE model by introducing a set of transit ISPs in between the content and eyeball ISPs. Again, the topology between any two connecting classes of ISPs is 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, the 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 forming 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 is $|\mathcal{N}| = |\mathcal{C}| + |\mathcal{T}| + |\mathcal{B}|$. Similarly, we define $\hat{\phi}_{T_k}$ as

the Shapley value revenue of T_k , and $\phi_{\mathcal{T}} = \sum_{T_k \in \mathcal{T}} \hat{\phi}_{T_k}$ as the aggregate Shapley value for the group of transit ISPs.

Theorem 2 (The Shapley Revenue for the CTE Model): We consider a network with a set \mathcal{C} of content ISPs, a set \mathcal{B} of eyeball ISPs, and a set \mathcal{T} of transit ISPs. 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 serve a single region. The Shapley value revenue of each ISP is in the following form:

$$\begin{aligned}\hat{\phi}_{B_j}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) &= \varphi_{\mathcal{B}}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) \hat{v}(\mathcal{N}) \quad \forall B_j \in \mathcal{B} \\ \hat{\phi}_{T_k}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) &= \varphi_{\mathcal{T}}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) \hat{v}(\mathcal{N}) \quad \forall T_k \in \mathcal{T} \\ \hat{\phi}_{C_i}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) &= \varphi_{\mathcal{C}}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) \hat{v}(\mathcal{N}) \quad \forall C_i \in \mathcal{C}\end{aligned}$$

where the normalized Shapley values $\varphi_{\mathcal{B}}$, $\varphi_{\mathcal{T}}$ and $\varphi_{\mathcal{C}}$ are

$$\begin{aligned}\varphi_{\mathcal{B}}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) &= \frac{1}{|\mathcal{N}|} \sum_{t=1}^{|\mathcal{T}|} \sum_{c=1}^{|\mathcal{C}|} \binom{|\mathcal{T}|}{t} \binom{|\mathcal{C}|}{c} \binom{|\mathcal{N}|-1}{t+c}^{-1} \\ \varphi_{\mathcal{T}}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) &= \frac{1}{|\mathcal{N}|} \sum_{b=1}^{|\mathcal{B}|} \sum_{c=1}^{|\mathcal{C}|} \binom{|\mathcal{B}|}{b} \binom{|\mathcal{C}|}{c} \binom{|\mathcal{N}|-1}{b+c}^{-1} \\ \varphi_{\mathcal{C}}(|\mathcal{B}|, |\mathcal{T}|, |\mathcal{C}|) &= \frac{1}{|\mathcal{N}|} \sum_{b=1}^{|\mathcal{B}|} \sum_{t=1}^{|\mathcal{T}|} \binom{|\mathcal{B}|}{b} \binom{|\mathcal{T}|}{t} \binom{|\mathcal{N}|-1}{b+t}^{-1}.\end{aligned}$$

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

$$\begin{aligned}\varphi_{\mathcal{B}}(b, t, c) &= \frac{1}{|\mathcal{N}|} \sum_{t'=1}^t \sum_{c'=1}^c \binom{t}{t'} \binom{c}{c'} \binom{b-1+t+c}{t'+c'}^{-1} \\ &= \varphi_{\mathcal{T}}(t, b, c) = \varphi_{\mathcal{C}}(c, t, b).\end{aligned}\quad (7)$$

To understand this symmetric property of the normalized Shapley value function, we can imagine that, as a group, the transit ISPs are as important as the content or the eyeball ISPs, because without the transit ISPs, the network is totally disconnected and cannot generate any revenue.

Fig. 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 $\hat{v}(\mathcal{N})$ to be 1. Along the x-axis, we vary the size of the content ISPs $|\mathcal{C}|$. For 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 $|\mathcal{B}|$ changes accordingly to satisfy $|\mathcal{B}| = |\mathcal{T}| - |\mathcal{C}|$. 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 the content and eyeball ISPs. We plot the value of $\phi_{\mathcal{T}}$ for $2 \leq |\mathcal{T}| \leq 8$. From Fig. 4, we can make two observations. First, similar to the result of the CE model, the ratio of $\phi_{\mathcal{C}} : \phi_{\mathcal{T}} : \phi_{\mathcal{B}}$ 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

$$\varphi_{\mathcal{B}}(b, t, c) = \xi \varphi_{\mathcal{B}}(\xi b, \xi t, \xi c) \quad \forall \xi = 1, 2, 3, \dots$$

For example, when we have only one ISP for each type, each ISP obtains one-third of the total revenue, i.e., $\varphi_{\mathcal{B}}(1, 1, 1) = 1/3$. The aggregate Shapley for each group of ISP will keep the




Fig. 4. Aggregate Shapley value revenue of the transit ISPs.

same, i.e., $\phi_B = \phi_T = \phi_C = \hat{v}(\mathcal{N})/3$, as long as the sizes of the groups of ISPs increase proportionally. If they do not increase proportionally, we have the second observation. Even as the number of transit ISPs keeps a constant as the sum of 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 with 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.

Observed by Labovitz *et al.* [14], new trends like consolidation and disintermediation, which interconnects contents to consumers directly, have been happening extensively in the last two years. Dhamdhere *et al.* [6] also confirmed the consolidation of the “core” of the Internet. Besides various reasons for that happening, e.g., disintermediation is driven by cost and performance, our Shapley revenue distribution result also rationalizes the economic incentives for ISPs to engage such activities so as to increase their Shapley revenues.

C. Multiple Contents and Regions Model

In this section, we extend our previous result for multiple contents and multiple regions. The conservation of revenue in (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 $CP_q^r = \beta_q \gamma_r^q X_r$ as the aggregate content-side revenue generated by providing content q for the users in region r . As a result, we can decompose the aggregate revenue $\hat{v}(\mathcal{N})$ as

$$\hat{v}(\mathcal{N}) = \sum_{r=1}^{|\mathcal{R}|} BP_r + \sum_{q=1}^{|\mathcal{Q}|} \sum_{r=1}^{|\mathcal{R}|} CP_q^r. \quad (8)$$

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 re-

gion r , these ISPs should get a fair share of the specific revenue component they contribute to generate.

Theorem 3 (Multiple Contents and Regions Model): We consider a network with a set \mathcal{C} of content ISPs, a set \mathcal{B} of eyeball ISPs, and a set \mathcal{T} of transit ISPs. Each content ISP C_i provides a set $Q_i \subseteq \mathcal{Q}$ of contents, and each eyeball ISP B_j serves a set $R_j \subseteq \mathcal{R}$ of regions. Both the content and eyeball ISPs are connected to the transit ISPs by a complete bipartite graph. The Shapley value revenue of each ISP is

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

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

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

where

$$h_q = \sum_{i=1}^{|\mathcal{C}|} \mathbf{1}_{\{q \in Q_i\}} \quad \forall q \in \mathcal{Q}$$

and

$$h_r = \sum_{j=1}^{|\mathcal{B}|} \mathbf{1}_{\{r \in R_j\}} \quad \forall r \in \mathcal{R}.$$

Theorem 3 shows that, in a multiple contents and regions environment, the Shapley value revenue can be expressed as separable Shapley components of specific content-side and eyeball-side revenues. h_r and h_q define the number of content ISPs that provide content q and the number of eyeball ISPs that serve region r , respectively. Notice that with ISP arrival or departure, these variables change accordingly. In particular, the eyeball-side revenue BP_r is not shared by the eyeball ISPs that are not serving 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 additivity property of the Shapley value [27], which can be applied to more general topologies of the network.

D. General Internet Topologies

Here, we consider more general network topologies than the complete bipartite connections assumed before. Because Tier-1 ISPs form a full mesh in practice, we focus on the topologies where the transit ISPs form a full mesh. However, our results apply for more general topologies.

To evaluate the Shapley value revenue distribution under a general topology, we first decompose the aggregate revenue according to (8). For each revenue component, i.e., BP_r or CP_q^r , a subsystem can be derived to distribute it. For example, the system in Fig. 3(d) can be decomposed into the six subsystems depicted in Fig. 5. To construct these subsystems, we only include the eyeball ISPs in region r for BP_r and the content ISPs that provide content q for CP_q^r . Because the topology is no

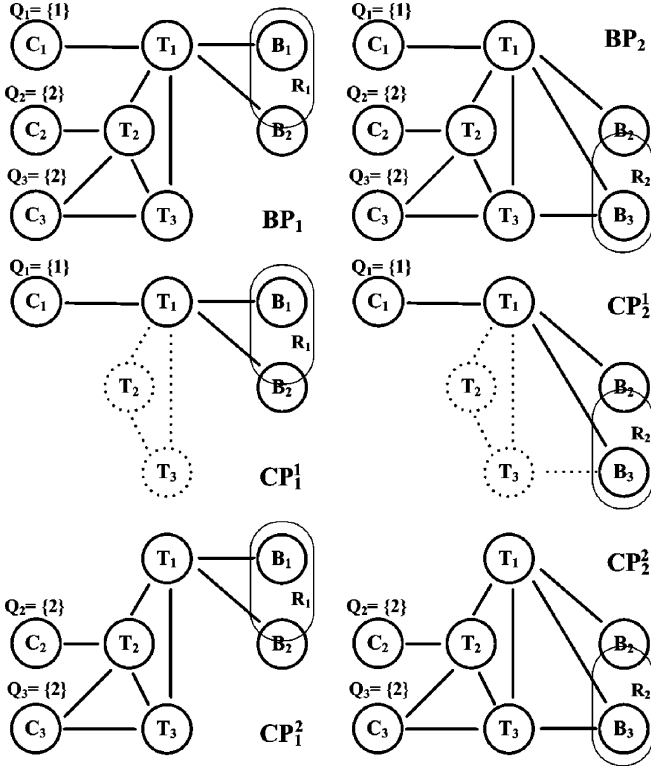


Fig. 5. Decomposition of the Shapley values.

longer a complete bipartite graph, some of the transit ISPs do not contribute for certain revenue components. We need to eliminate these *dummy* transit ISPs as depicted in dotted circles in Fig. 5.

Definition 4: An ISP i is *dummy* with respect to a worth function v if $\Delta_i(v, \mathcal{S}) = 0$ for every $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$.

The remaining problem is to derive the Shapley revenue distribution for each of the decomposed subsystems in Fig. 5. Theorem 2 gives the closed-form solution for complete bipartite topologies. Although topological changes affect the Shapley value revenues, we can evaluate them via a dynamic programming procedure if the system is *canonical*.

Definition 5: 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 (\mathcal{N}, \hat{v}) , the aggregate revenue $\hat{v}(\mathcal{N})$ is either wholly earned or wholly 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 Fig. 2. Otherwise, only partial revenue will be lost, i.e., $0 < \hat{v}(\mathcal{S}) < \hat{v}(\mathcal{N})$, which violates the canonical property. By having a full-mesh among the transit ISPs and elastic intraregion user demands, each of the decomposed subsystems is indeed a *canonical* system. Thus, any coalition \mathcal{S} obtains either the whole decomposed revenue or nothing.

Definition 6: An ISP i is called a *veto ISP* with respect to worth function v if i belongs to all \mathcal{S} with $v(\mathcal{S}) > 0$.

Every veto ISP with respect to \hat{v} 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 users. Now, we are ready to describe the dy-

amic programming procedure that derives 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 $\varphi_i(\mathcal{N}, v)$ 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}, \hat{v}) happens to have a complete bipartite topology.

E. Generalized Demand Elasticity

In the previous sections, we showed that in a general topology with multiple regions and contents, we can decompose the aggregate revenue into eyeball-side components BP_r and content-side components CP_r^q . Each revenue component can be distributed to ISPs in a canonical system. This is a result based on the elastic intraregion user demand assumption.

Here, we relax the intraregion demand assumption. Without loss of generality, we focus on an eyeball-side revenue component $BP_r = \alpha_r X_r$ in region r . For each eyeball ISP B_j , we denote e_j as the percentage of users x_j^r whose demands are elastic to other ISPs in the same region. Thus, $1 - e_j$ represents the percentage of *inelastic* users that will leave the system when ISP j becomes unavailable in region r . The elasticity parameter e_j can be used to model differentiated services provided by the ISPs, such that a fixed number of users is sticky to a certain ISP for its special service.

Because certain users are sticky to the eyeball ISPs, when an eyeball ISP B_j leaves the system, the system will lose $\alpha_r(1 - e_j)x_j^r$ amount of revenue, and therefore the system is not canonical. However, this generalized user demand can also be decomposed into canonical systems where the dynamic programming procedure can be used to compute the Shapley values. The idea is to separate the proportion of inelastic user demands as if they are from a different region

$$BP_r = \alpha_r X_r = \alpha_r \sum_{r \in R_j} e_j x_j^r + \sum_{r \in R_j} \alpha_r (1 - e_j) x_j^r. \quad (9)$$

In (9), the first term is the aggregate revenue generated by all elastic users and would be shared in a canonical system of all ISPs, i.e., $\mathcal{N} = \mathcal{C} \cup \mathcal{T} \cup \mathcal{B}$. Each of the remaining terms rep-

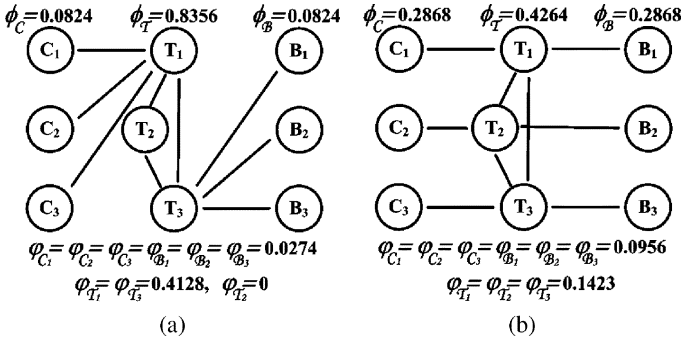


Fig. 6. Two extreme ways to interconnect. (a) Focused connection. (b) Uniform connection.

represents the revenue generated by the inelastic users of a certain ISP B_j and would be shared in a canonical system of ISPs $C \cup T \cup \{B_j\}$ because other eyeball ISPs are *dummy* with respect to the inelastic users of B_j . From (9), we know that the Shapley revenue of an eyeball ISP B_j is proportional to its inelastic user base $(1 - e_j)x_j^r$ and the aggregate elastic user base $\sum_{r \in R_j} e_j x_j^r$. However, the weight on $(1 - e_j)x_j^r$ is higher (because other eyeball ISPs are dummy to this revenue component). This also explains why ISPs of the same type obtain the same amount of the Shapley revenue (Theorem 1 and 2) regardless of the user base x_j^r , when the intraregion user demand is fully elastic.

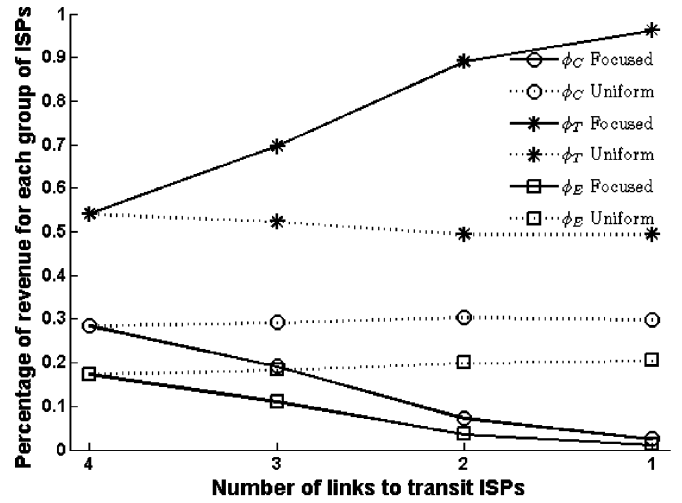
As a result, a nonelastic intraregion demand can be decomposed into canonical subsystems where demands are elastic.

F. Connectivity Effects on the Shapley Revenues

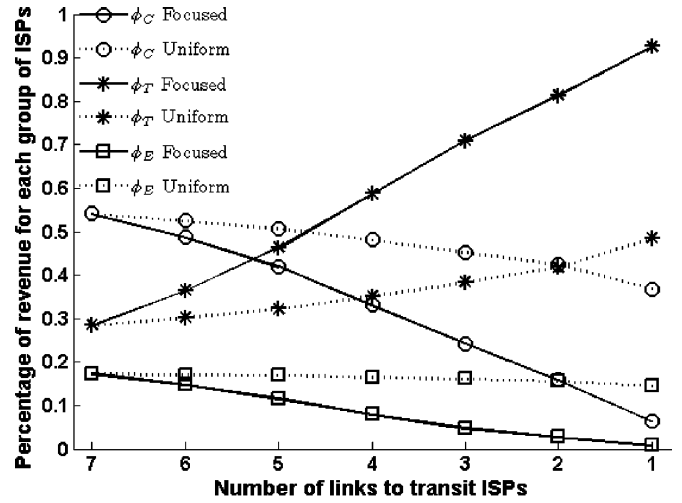
For each canonical system, the Shapley value distribution only depends on the topological structure of the ISPs. Here, we explore how the topological structure affects the Shapley value distribution on a canonical system. Fig. 4 compares the aggregate Shapley revenue for each group of ISPs when the number of ISPs in each group changes. 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.

Fig. 6 illustrates the Shapley revenue distribution for ISPs when $|C| = |T| = |B| = 3$ and each content or eyeball ISP only connects to one transit ISP. We can see that the Shapley revenues differ drastically, depending on how the content and eyeball ISPs are connected to the transit ISPs. Fig. 6(a) shows the case where 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. Fig. 6(b) shows the case where all content and eyeball ISPs are connected to transit ISPs uniformly. In contrast, the group of transit ISPs only obtains 42%, half of the previous share, of the total revenue.


Fig. 7 illustrates how the values of ϕ_B , ϕ_T , and ϕ_C vary reacting to the changes of the interconnecting links. Along the x-axis, we vary the degree of connectivity to the transit ISPs. We start from the complete bipartite topology where each content or eyeball ISP connects to all $|T|$ transit ISPs. Then, we gradually decrease the number of transit ISPs to which each content or eyeball ISP connects. We plot two types of topologies.



(a)



(b)



(c)

Fig. 7. Shapley value revenue for the groups of ISPs. (a) ISP groups with $|C| = 7$, $|T| = 4$, $|B| = 10$. (b) ISP groups with $|C| = 4$, $|T| = 7$, $|B| = 10$. (c) ISP groups with $|C| = 4$, $|T| = 10$, $|B| = 7$.

- 1) Focused connections (solid lines): For any degree of connectivity κ , the content ISPs connect to the first κ transit

ISPs and the eyeball ISPs connect to the last κ transit ISPs [as in Fig. 6(a)].

- 2) Uniform connections (dotted lines): The content and eyeball ISPs connect to the transit ISPs in a round-robin manner, where each transit ISP connects to approximately the same number of content and eyeball ISPs [as in Fig. 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 $\kappa = 1$, 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 the transit ISPs. In an extreme case of Fig. 7(a) where $|\mathcal{T}|$ is small relative to $|\mathcal{C}|$ and $|\mathcal{B}|$, the value of ϕ_T might decrease a little bit when the 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 the content and eyeball ISPs connect to the transit ISPs strongly affect the value ϕ_T .

V. SHAPLEY COST DISTRIBUTION

In this section, we explore the Shapley cost distribution for ISPs. Parallel to the previous section, we first derive the Shapley cost under the CE and the CTE models.

Theorem 5 (The Shapley Cost for the CE Model): We consider a set \mathcal{C} of content ISPs and a set \mathcal{B} of eyeball ISPs, under the CE model with a complete bipartite graph topology. The Shapley value cost for each ISP is

$$\begin{aligned}\bar{\varphi}_{B_j}(|\mathcal{B}|, |\mathcal{C}|) &= \frac{|\mathcal{C}|BL_j}{|\mathcal{C}|+1} + \frac{1}{|\mathcal{B}|(|\mathcal{B}|+1)} \sum_{C_i \in \mathcal{C}} CL_i \quad \forall B_j \in \mathcal{B} \\ \bar{\varphi}_{C_i}(|\mathcal{B}|, |\mathcal{C}|) &= \frac{|\mathcal{B}|CL_i}{|\mathcal{B}|+1} + \frac{1}{|\mathcal{C}|(|\mathcal{C}|+1)} \sum_{B_j \in \mathcal{B}} BL_j \quad \forall C_i \in \mathcal{C}.\end{aligned}$$

Theorem 6 (The Shapley Cost for the CTE Model): We consider a network with a set \mathcal{C} of content ISPs, a set \mathcal{B} of eyeball ISPs, and a set \mathcal{T} of transit ISPs. Both the content and the 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 serve a single region. The Shapley value revenue for each ISP is

$$\begin{aligned}\bar{\varphi}_{B_j} &= \varphi_B(|\mathcal{B}|, 1, |\mathcal{C}|) \sum_{T_k \in \mathcal{T}} TL_k + \varphi_B(|\mathcal{B}|, |\mathcal{T}|, 1) \sum_{C_i \in \mathcal{C}} CL_i \\ &\quad + \varphi_B(1, |\mathcal{T}|, |\mathcal{C}|)BL_j \quad \forall B_j \in \mathcal{B} \\ \bar{\varphi}_{T_k} &= \varphi_T(1, |\mathcal{T}|, |\mathcal{C}|) \sum_{B_j \in \mathcal{B}} BL_j + \varphi_T(|\mathcal{B}|, |\mathcal{T}|, 1) \sum_{C_i \in \mathcal{C}} CL_i \\ &\quad + \varphi_T(|\mathcal{B}|, 1, |\mathcal{C}|)TL_k \quad \forall T_k \in \mathcal{T}\end{aligned}$$

$$\begin{aligned}\bar{\varphi}_{C_i} &= \varphi_C(1, |\mathcal{T}|, |\mathcal{C}|) \sum_{B_j \in \mathcal{B}} BL_j + \varphi_C(|\mathcal{B}|, 1, |\mathcal{C}|) \sum_{T_k \in \mathcal{T}} TL_k \\ &\quad + \varphi_C(|\mathcal{B}|, |\mathcal{T}|, 1)CL_i \quad \forall C_i \in \mathcal{C}\end{aligned}$$

where φ_B , φ_T , and φ_C are the normalized Shapley value function defined in (7).

From Theorems 5 and 6, we can make two observations: 1) the Shapley cost of each ISP is a separable function on individual cost components BL_j , TL_k , and CL_i ; and 2) each cost component is shared by ISPs under the same distribution functions in Theorem 1 and 2, respectively, assuming all other ISPs of the same type as the cost originator are dummy. Analogously, we can imagine that the cost of an ISP is *inelastic* to the ISPs of the same type. Consequently, under general topologies, we can separate each individual cost component as a canonical system and calculate the Shapley value of that cost component among the ISPs by the dynamic programming procedure developed in Theorem 4.

VI. IMPLICATIONS AND JUSTIFICATIONS

In the previous two sections, we developed the Shapley revenue $\hat{\varphi}$ and cost $\bar{\varphi}$ distribution for ISPs under general Internet topologies. To achieve a Shapley profit distribution, each ISP, for example an eyeball ISP B_j , should receive a payment of $\varphi_{B_j} + BL_j$ that guarantees a profit of φ_{B_j} after covering its cost BL_j . We can rewrite the payment as

$$\varphi_{B_j} + BL_j = \hat{\varphi}_{B_j} - \bar{\varphi}_{B_j} + BL_j = \hat{\varphi}_{B_j} + (BL_j - \bar{\varphi}_{B_j}). \quad (10)$$

The above equation tells that we can first make a Shapley revenue distribution $\hat{\varphi}$ to all ISPs, and then make a cost adjustment for each ISP, e.g., pay $BL_j - \bar{\varphi}_{B_j}$ to ISP B_j . Based on Theorem 6, we know that if the cost BL_j is relatively higher than the cost of TL_k 's and CL_i 's, the cost adjustment will be large, and vice versa. Although the individual cost, e.g., BL_j , will be recovered, the Shapley cost component, e.g., $\bar{\varphi}_{B_j}$, is a function of the individual cost, and therefore ISPs will have the incentives to reduce their individual costs so as to maximize their profits. Detailed models and results on general incentives and routing costs can be found in [17].

Although the Shapley value solution inherits multiple desirable properties, the actual profit distribution in the Internet might be different from the Shapley value due to the inefficient bilateral agreements between the 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 structures for differentiated services. We start with a solution concept, *the core* [28], which leads to a brief discussion of the stability of the Shapley value solution.

A. Core, Convex Game, and Stability

Analogous to the concept of Nash equilibrium in noncooperative games, the core is a stable solution concept for coalition games where no deviation from the grand coalition \mathcal{N} will be profitable. Let a vector ϕ be a solution of a coalition game, where each ϕ_i is the profit shared by player i . We confine ourselves to the set of *feasible* solutions that share the value of $v(\mathcal{N})$ among all players, i.e., $\sum_{i \in \mathcal{N}} \phi_i = v(\mathcal{N})$.

Definition 7: The *core* of a coalition game (\mathcal{N}, v) is the set of feasible solutions ϕ that satisfies

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

The efficiency property of the Shapley value [27] makes it a feasible solution. Under our profit distribution context, stability concerns whether ISPs can form a coalition to earn more profit than the aggregate Shapley profit. If so, ISPs do not have incentives to cooperate all together and may deviate from the Shapley profit distribution. Mathematically, it requires the Shapley value to be in the core

$$\sum_{i \in \mathcal{S}} \varphi_i(\mathcal{N}, v) \geq v(\mathcal{S}) \quad \forall \mathcal{S} \subseteq \mathcal{N}. \quad (11)$$

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. In general, the core of a coalition game might not even exist, which means no feasible solution can be stable. Fortunately, it has been shown that many types of coalition games have a nonempty core. One of them is the *convex game*.

Definition 8: A coalition game (\mathcal{N}, v) is *convex* if the worth function v is convex, i.e., for all coalition \mathcal{S} and \mathcal{S}' , v satisfies

$$v(\mathcal{S}) + v(\mathcal{S}') \leq v(\mathcal{S} \cup \mathcal{S}') + v(\mathcal{S} \cap \mathcal{S}'). \quad (12)$$

Notice that the convexity condition is pretty loose. Naturally, players will cooperate to form a bigger coalition if they can achieve a higher aggregate value than the sum of individual values. Mathematically, this can be described as a *super-additive* condition, i.e., $v(\mathcal{S}) + v(\mathcal{S}') \leq v(\mathcal{S} \cup \mathcal{S}')$. The convexity condition is implied by satisfying the super-additive condition. On the other hand, if the super-additive condition cannot be satisfied, players generally do not have incentive to cooperate, and therefore the core might be empty as well.

The Shapley value is known to be in the core [28] of a convex game. Particularly, Lloyd Shapley proved that the marginal contributions $\Delta_i(v, \mathcal{S}(\pi, i))$ defined in (2) form the vertices of the core of the convex game. The Shapley value, which is the average of vertices of the core, is located at the center of gravity of the core. This result shows that besides being a stable solution, the Shapley value is also the most *robust* solution among all stable solutions. Fig. 8 illustrates the core (solid line segment) of a two-ISP example. The x-axis and y-axis represent the profit 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 at the midpoint 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 .

From a practical point of view, to implement the Shapley value solution, ISPs need to divulge topological information as well as their cost structures, which ISPs do not want to reveal. Therefore, a centralized authority might be needed to enforce the process. However, without a Shapley value mechanism, ISPs might still reach a stable solution in the core. Our vision is that the Shapley mechanism can be used whenever ISP disputes




Fig. 8. Core of a two-ISP example.

happen and government/regulatory forces are needed to restabilize the interconnections and settlements.

B. Justifications for Stable Bilateral Agreements

The Shapley value solution suggests a value chain illustrated in Fig. 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 plus the cost adjustment of the group, of the payments and forwards the remaining along the network. However, in practice, ISPs negotiate bilateral settlements. Huston [11] concluded that the zero-dollar peering and the customer/provider relationships were the only stable models for the Internet at the '90s. The effective profit distribution resulted from these bilateral agreements should probably be different from that of the Shapley value distribution. One may wonder why these bilateral agreements were stable and how close they were to the Shapley solution. Fig. 10 illustrates the scenario when local ISPs were still homogeneous and the end-to-end traffic patterns exhibited symmetry at the '90s. At that time, local ISPs were not specialized to be content or eyeball ISPs. They obtain end-payments from content providers and/or residential users. To achieve the Shapley profit distribution, ISPs need to exchange different payments. Each local ISP forwards more money to the transit ISPs than they receive from them. Effectively, the net money exchange would be from the local ISPs to the transit ISPs. Due to the symmetric traffic pattern, the net money exchange between the transit ISPs would be close to zero. This result coincides with the zero-dollar peering and the customer/provider relationships established from bilateral agreements. Although the exact profit distribution might still be different from the Shapley value, we conjecture that the resulting profit distribution was very close to the Shapley value solution so that it was in the core, and thus stable.

C. Justifications for Unstable Bilateral Agreements

The Internet has been changing dramatically during the past two decades. Traditional content providers have developed multibillion businesses from the Internet via advertising (e.g., Google), e-commerce (e.g., Amazon and Ebay), and other services (e.g., Yahoo! and Bloomberg). They also build infrastructures like cloud computing platforms (e.g., Amazon EC2 and Google App Engine) and behave more like content ISPs. On the other hand, many traditional transit ISPs nowadays provide Internet access to millions of end-users and behave more like eyeball ISPs. These changes affect the ISPs in two ways. First, the revenue flows become very different: Content providers make a large amount of revenue by providing applications over the Internet. However, due to the flat-rate pricing scheme for end-users, the per-user revenue earned from the eyeball side

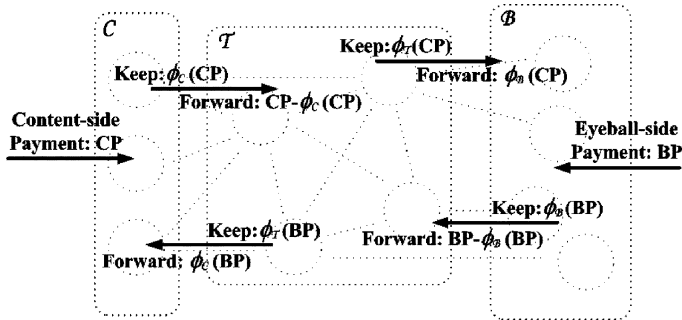


Fig. 9. Value chain to implement the Shapley profit.

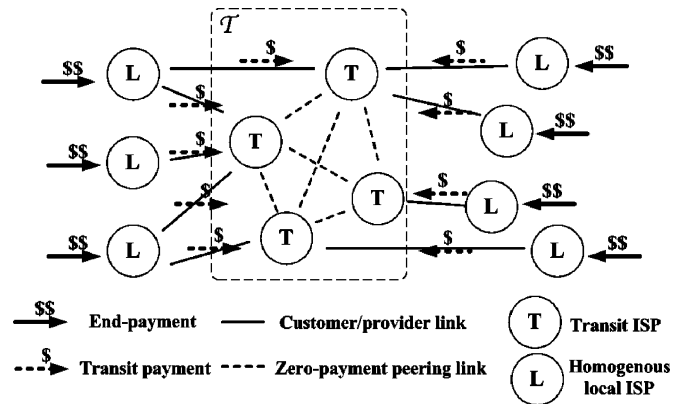


Fig. 10. Traditional ISP structure with homogeneous local ISPs and transit ISPs.

does not change much. Second, due to the pervasive use of P2P technologies and multimedia applications like video streaming and voice over IP, the traffic volume and the corresponding routing costs that the eyeball and transit ISPs have to bear have increased dramatically. These trends are shown by the Internet observatory project [14] where researchers found that the price of wholesale bandwidth decreases while the growth of advertising contents keeps increasing. Consequently, the Shapley value solution in the new environment should change as well.

Faratin *et al.* [7] 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*.¹ Paid-peering is identical to zero-dollar peering in terms of traffic forwarding, except that one party needs to pay another. Because zero-dollar peering at the Tier-1 level often requires participating ISPs to be *transit-free*, paid-peering makes it possible for some very large ISPs to satisfy the letter of this requirement before they achieve the coveted Tier-1 status. By applying the Shapley profit distribution to the content–transit–eyeball model, we justify and fortify the rationale of paid-peering between transit ISPs.

Fig. 11 illustrates a scenario where the content ISPs connect to a set of content-side transit (CT) ISPs and the eyeball ISPs connect to a set of eyeball-side transit (BT) ISPs. The eyeball-side revenues are much smaller than the content-side revenues because *BP*'s are based on fixed monthly payments from

¹Despite being uncommon in the early days, paid-peering might have been around since the late 1990s according to anecdotal evidence from operators.

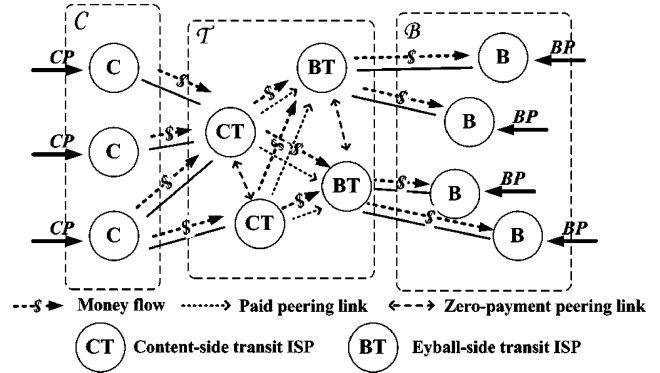


Fig. 11. Shapley value implied money exchange.

residential users and *CP*'s are growing with the Internet-related businesses. After netting the exchange of payments, including the cost adjustments, along the value chain in Fig. 9, we show the resulting bilateral money flows that implement the Shapley value solution in Fig. 11. We observe that the content ISPs obtain content-side revenues and pay the CT ISPs. This is the same customer/provider relationship as before. However, the zero-dollar peering relationship does not happen between all pairs of the transit ISPs. Notice that the CT ISPs need to forward the content-side value toward the eyeball-side, which creates the paid-peering relationship emerged with heterogeneous ISPs.

One unconventional observation is that the eyeball ISPs need to receive compensations from the content side through the BT ISPs. This implies that the transit ISPs should pay the 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 profit-sharing solution might be located outside the core, and thus becomes 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. However, whether we can achieve this largely depends on the willingness of profit-sharing of the content providers, e.g., Google and Amazon. Consequently, incumbent and transit ISPs want to create service differentiations so as to generate extra profits. This naturally leads to the debate of *network neutrality*.

D. 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 Internet. The proponents [4], [30] criticized the discriminatory behaviors of the ISPs, believing that they harm the productivity, innovation, and end-to-end connectivity of the Internet. However, the opponents [13] advocated that offering premium services would stimulate innovations on the edges of the network. Musacchio *et al.* [22] showed that different parameters, e.g., advertising rate and end-user price sensitivity, influence whether a neutral or nonneutral regime achieves a higher social welfare.

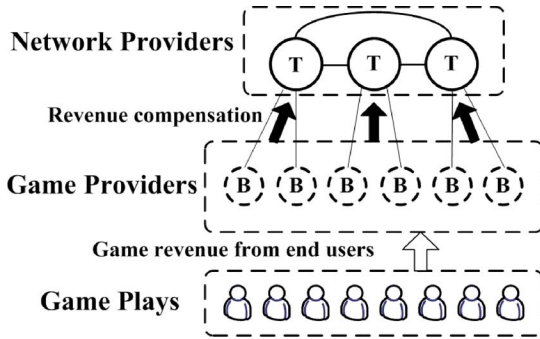


Fig. 12. Compensation structure for game services.

As we discussed before, bilateral agreements that severely deviate from the Shapley profit 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 profit distribution mechanism, ISPs do not have the incentive to architect the cooperative provisioning for such services. As a generic profit-sharing mechanism, the Shapley value solution can also be used to encourage ISPs to participate and fairly share profits. Here, we illustrate two potential differentiated services and the implied compensation structures for the supporting ISPs.

1) *Supporting Gaming Services:* The booming online gaming industry has brought huge profits to game providers. 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 guarantee service qualities so as to support game applications.

Fig. 12 illustrates the compensation structure implied by the Shapley value solution for the game services. By providing gaming applications to players, the 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, who provide the interconnections between the game players and different game providers. Due to the symmetric traffic patterns of network games, the compensation structure is similar to the customer/provider and the zero-dollar peering structure in Fig. 10. In Fig. 12, the game providers need to compensate network providers for supporting the new service and the network providers connect to one another with zero-dollar peering agreements.

2) *Supporting Real-Time Data Services:* Another potential application is real-time data services across the Internet. Many real-time applications require a low latency to retrieve accurate real-time data, e.g., stock/option quotes and sports game scores; others require guaranteed network services, e.g., trading transactions and online gamble. High (or asymmetric) latencies can make these applications extremely vulnerable, as a few milliseconds here or there can translate to billions of dollars with auto-

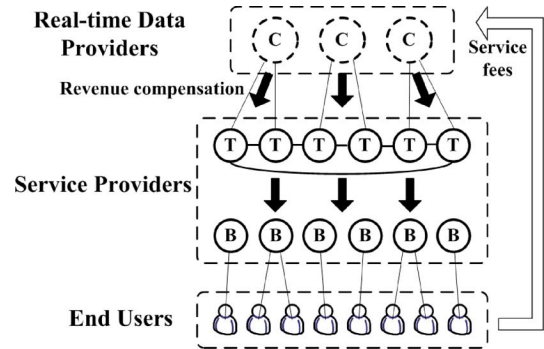


Fig. 13. Implied compensation structure for real-time data services.

mated trading. Network security is another major concern for providing secured transactions over the network. In order to implement more robust and secure protocols across the Internet, cooperation among ISPs might be needed to support low latency and to prevent malicious attackers and information stealers.

Fig. 13 illustrates the compensation structure implied by the Shapley value solution for secure real-time data services. On the top, the data providers (dotted circles) earn revenue from customers by providing the online data services. They can be considered as the content providers who obtain revenue by attracting business from their customers’ ability to access the Internet. Service providers (solid circles) cooperatively implement secure protocols to support the interactions between the data providers and their customers. The ISPs directly connected to the customers are like eyeball ISPs, and the intermediate ISPs are like transit ISPs. Due to the asymmetric traffic characteristic, the compensation structure is similar to the one shown in Fig. 11. In Fig. 13, the data providers compensate the transit ISPs the same as in a customer/provider relationship. The transit ISPs need to compensate the eyeball ISPs the same as in a reverse customer/provider relationship. Paid-peering relationship might also exist if there are multiple levels of transit ISPs in the network.

VII. RELATED WORK

Our previous work [17] proposed a clean-slate Shapley profit distribution mechanism for ISPs 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 the individual profits and the system’s social welfare. Due to the multilateral 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 [15] to model a detailed Internet structure was limited to the content–eyeball model introduced by Faratin *et al.* [7]. In this paper, we extend our model to include a third class of ISPs: the transit ISPs. We generalize all results in [15] 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 [2] and Huston [11] started exploring the interconnection settlements of the ISP in the 1990s. Huston [11] and Frieden compared the existing Internet settlement models to 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 [7] exhibits interconnection disputes in the Internet and observes the emergence of a *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 the rationale of a paid-peering relationship between transit ISPs. Moreover, it also suggests that a reverse 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 [10] 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.* [3] experimented on AS relationships and observed violations of the valley-free property [10] from BGP routing tables. Our work treats ISPs as cooperative entities that form coalitions to share profit. 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 [4], [9], [30] started when discriminatory practices, e.g., selectively dropping packets, were found with broadband provider and cable operators. Crowcroft [4] reviewed technical aspects of network neutrality and concluded that network neutrality should not be engineered. Both sides of the debate are concerned about whether differentiated services should be provided in the Internet. Musacchio *et al.* [22] 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 depends 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 [21], game theory [23] has been used to address pricing [26] and incentive problems [19] in networking areas. Unlike the majority of noncooperative game models, the Shapley value [24] originates from *coalition games* [23] that model the cooperative nature of groups. Eyal Winter's survey [29] provides a thorough investigation on the Shapley value and its properties.

VIII. CONCLUSION

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 evaluate 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 profit 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 of the symmetry of the traffic flows and the homogeneity of the 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 solution resulting from bilateral agreements severely deviates from the Shapley value solution, which implies a paid-peering relationship between the transit ISPs and a reverse customer/provider relationship between the transit and the eyeball ISPs. We conjecture that many of the failures of the existing agreements are due to the lack of implementing these paid-peering and reverse customer/provider relationships via bilateral agreements. 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 the newly brought revenues and to enrich the Internet services. We believe that our results can be useful for ISPs to settle bilateral disputes and for regulatory institutions to regulate the Internet industry.

REFERENCES

- [1] Y. Bachrach, E. Markakis, A. D. Procaccia, J. S. Rosenschein, and A. Saberi, "Approximating power indices," in *Proc. 7th AAMAS*, Estoril, Portugal, 2008, pp. 943–950.
- [2] J. P. Bailey, "The economics of Internet interconnection agreements," in *Internet Economics*, L. W. McKnight and J. P. Bailey, Eds. Cambridge, MA: MIT Press, 1997.
- [3] 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 Trans. Netw.*, vol. 15, no. 2, pp. 267–280, Apr. 2007.
- [4] J. Crowcroft, "Net neutrality: The technical side of the debate: A white paper," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 1, pp. 49–56, Jan. 2007.
- [5] G. Demange and M. Wooders, *Group Formation in Economics: Networks, Clubs, and Coalitions*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [6] A. Dhamdhere and C. Dovrolis, "Ten years in the evolution of the Internet ecosystem," in *Proc. 8th ACM SIGCOMM IMC*, Vouliagmeni, Greece, Oct. 2008, pp. 183–196.
- [7] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr, "Complexity of Internet interconnections: Technology, incentives and implications for policy," in *Proc. 35th TPRC*, Arlington, VA, Sep. 2007.
- [8] R. Frieden, "Without public peer: The potential regulatory and universal service consequences of Internet Balkanization," *Virginia J. Law Technol.*, vol. 3, no. 8, pp. 1522–1687, 1998.
- [9] R. Frieden, "Network neutrality or bias? Handicapping the odds for a tiered and branded Internet," Working Paper 1755, 2006, bepress Legal Ser.
- [10] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, Dec. 2001.
- [11] G. Huston, *ISP Survival Guide: Strategies for Running a Competitive ISP*. New York: Wiley, 1999.
- [12] M. O. Jackson, "Allocation rules for network games," *Games Econ. Behav.*, vol. 51, no. 1, pp. 128–154, Apr. 2005.
- [13] M. Jamison and J. Hauge, "Getting what you pay for: Analyzing the net neutrality debate," Dept. Econ., Univ. Florida, Gainesville, PURC Working Paper, 2008 [Online]. Available: <http://ssrn.com/abstract=1081690>
- [14] C. Labovitz, D. McPherson, and S. Iekel-Johnson, "Internet observatory 2009 annual report," in *Proc. NANOG47*, Oct. 2009.

[15] 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," in *Proc. ACM NetEcon*, Seattle, WA, Aug. 2008, pp. 61–66.

[16] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "On cooperative settlement between content, transit and eyeball Internet service providers," in *Proc. 2008 ACM CoNEXT*, Madrid, Spain, Dec. 2008, Article No. 7.

[17] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "Internet economics: The use of Shapley value for ISP settlement," *IEEE/ACM Trans. Netw.*, vol. 18, no. 3, pp. 775–787, Jun. 2010.

[18] R. T. B. Ma, D. M. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "The Shapley profit for content, transit and eyeball ISPs," EE Dept., Columbia Univ., New York, Tech. Rep., 2010 [Online]. Available: <http://dna-pubs.cs.columbia.edu/citation/paperfile/169/TR.pdf>

[19] 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 Trans. Netw.*, vol. 14, no. 5, pp. 978–991, Oct. 2006.

[20] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, K. C. Claffy, and A. Vahdat, "Lessons from three views of the Internet topology," Tech. rep. CAIDA-TR-2005-02, 2005.

[21] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford, U.K.: Oxford Univ. Press, 1995.

[22] J. Musacchio, G. Schwartz, and J. Walrand, "Network neutrality and provider investment incentives," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Nov. 2007, pp. 1437–1444.

[23] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA: MIT Press, 1994.

[24] A. Roth, *The Shapley Value: Essays in Honor of Lloyd S. Shapley*. Cambridge, U.K.: Cambridge Univ. Press, 1988.

[25] T. Roughgarden and E. Tardos, "How bad is selfish routing?," *J. ACM*, vol. 49, no. 2, pp. 236–259, Mar. 2002.


[26] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple ISPs," *IEEE/ACM Trans. Netw.*, vol. 14, no. 6, pp. 1233–1245, Dec. 2006.

[27] L. Shapley, "A value for n -person games," in *Contributions to the Theory of Games II*, H. W. Kuhn and A. W. Tucker, Eds. Princeton, NJ: Princeton Univ. Press, 1953, vol. 28, *Annals of Mathematics Studies*, pp. 307–317.

[28] L. Shapley, "Cores of convex games," *Int. J. Game Theory*, vol. 1, no. 1, pp. 11–26, Dec. 1971.

[29] E. Winter, "The Shapley value," in *The Handbook of Game Theory*, R. J. Aumann and S. Hart, Eds. Amsterdam, The Netherlands: North-Holland, 2002.

[30] T. Wu, "Network neutrality, broadband discrimination," *J. Telecommun. High Technol. Law*, vol. 2, pp. 141–175, 2003.



Richard T. B. Ma received the B.Sc. degree (first-class honors) in computer science and the M.Phil. degree in computer science and engineering from the Chinese University of Hong Kong in July 2002 and July 2004, respectively, and the Ph.D. degree in electrical engineering from Columbia University, New York, NY, in May 2010.


During his Ph.D. study, he worked as a Research Intern with IBM T. J. Watson Research Center, Hawthorne, NY, and Telefonica Research, Barcelona, Spain. He is currently a Research Scientist with the Advanced Digital Sciences Center, University of Illinois at Urbana-Champaign, and an Assistant Professor with the School of Computing, National University of Singapore. His current research interests include distributed systems, network economics, game theory, and stochastic processes.



Dah Ming Chiu (SM'02–F'08) received the B.Sc. degree in electrical engineering from Imperial College London, London, U.K., in 1975, and the Ph.D. degree in applied mathematics from Harvard University, Cambridge, MA, in 1980.

After 20 years in industry (Bell Labs, DEC, and Sun), he is now a Professor with the Department of Information Engineering, The Chinese University of Hong Kong. His research interests include network economics and architecture, P2P networking, and wireless networks.

Prof. Chiu is an IEEE Fellow recognized for his "contribution to distributed resource allocation algorithms in computer networks." He serves as an Associate Editor for the IEEE/ACM TRANSACTION ON NETWORKING and is regularly invited to serve on the Technical Program Committees of IEEE ICNP, IEEE INFOCOM, ACM CoNext, IEEE IWQoS, and recently ACM SIGCOMM as well.



John C. S. Lui (M'93–SM'02–F'10) was born in Hong Kong. He received the Ph.D. degree in computer science from the University of California, Los Angeles, in 1992.

He is currently the Chairman of the Department of Computer Science and Engineering, The Chinese University of Hong Kong (CUHK). He was with the IBM San Jose/Almaden Lab, San Jose, CA, after graduation and later joined the CUHK. His current research interests are in theoretic/applied topics in data networks, distributed multimedia systems, network security, and mathematical optimization and performance evaluation theory.

Dr. Lui served as the Technical Program Committee Co-Chair of ACM SIGMETRICS 2005 and General Co-Chair of the IEEE International Conference on Network Protocols (ICNP) 2007, and is currently the Vice-President of ACM SIGMETRICS. He received various departmental teaching awards and the CUHK Vice-Chancellor's Exemplary Teaching Award in 2001. He is a co-recipient of Best Paper Award in IFIP WG 7.3 Performance 2005 and the IEEE/IFIP Network Operations and Management (NOMS) Conference.

Dr. Lui served as the Technical Program Committee Co-Chair of ACM SIGMETRICS 2005 and General Co-Chair of the IEEE International Conference on Network Protocols (ICNP) 2007, and is currently the Vice-President of ACM SIGMETRICS. He received various departmental teaching awards and the CUHK Vice-Chancellor's Exemplary Teaching Award in 2001. He is a co-recipient of Best Paper Award in IFIP WG 7.3 Performance 2005 and the IEEE/IFIP Network Operations and Management (NOMS) Conference.



Vishal Misra (S'98–M'99) received the B.Tech. degree from the Indian Institute of Technology, Bombay, India, in 1992, and the Ph.D. degree from the University of Massachusetts Amherst, in 2000, both in electrical engineering.

He is an Associate Professor of computer science with Columbia University, New York, NY. His research emphasis is on mathematical modeling of computer systems, bridging the gap between practice and analysis. His recent work includes the areas of Internet economics, peer-to-peer networks, and efficient scheduling policies.

Dr. Misra served as the Guest Editor for the *Journal of Performance Evaluation* and the Technical Program Committee Chair and General Chair of the ACM SIGMETRICS conference. He has participated as a member of program committees for conferences such as the IEEE INFOCOM, ACM SIGMETRICS, ACM SIGCOMM, IFIP Performance, and IEEE ICNP. He serves on the Board of Directors of ACM SIGMETRICS. He has received an NSF CAREER Award, a DoE CAREER Award, and an IBM Faculty Awards.



Dan Rubenstein (M'00) received the B.S. degree in mathematics from the Massachusetts Institute of Technology, Cambridge, in 1992, the M.A. degree in math from the University of California, Los Angeles, in 1994, and the Ph.D. degree in computer science from the University of Massachusetts Amherst in 2000.

He is now an Associate Professor of electrical engineering and computer science with Columbia University, New York, NY. He is generally interested in the design and performance of networked systems

and also has interests in the security and integrity of such systems as well as the general design of distributed communication algorithms.

Dr. Rubenstein has received a U.S. National Science Foundation CAREER Award, an IBM Faculty Award, an IEEE ICNP Best Paper Award, and (as a student) an ACM SIGMETRICS Best Student Paper Award.