

A Preference Model for Deciding the Market Share of Network Service Providers

Richard T. B. Ma
School of Computing
National University of Singapore
Email: tbma@comp.nus.edu.sg

John C. S. Lui
Dept. of Computer Science & Engineering
The Chinese University of Hong Kong
Email: cslui@cse.cuhk.edu.hk

Vishal Misra
Dept. of Computer Science
Columbia University
misra@cs.columbia.edu

Abstract—Started with a single *best-effort* service, the Internet has evolved to an ecosystem where different Service Providers (SPs), e.g., Internet Service Providers (ISPs) and Content Distribution Networks (CDNs) provide different types of services, e.g., IP transit and content caching/distribution. In this paper, we propose a preference model on how the Content Providers (CPs) of the Internet choose SPs based on their own characteristics as well as the quality and price of the SPs. Based on the preferences of the CPs and the available set of SPs, our model predicts the collective choices of the CPs and the resulting market share of the SPs. Our model provides a better understanding of the business relationship among the CPs and the SPs, and can help them to make informed decisions so to succeed in this competitive Internet ecosystem. Furthermore, it can also help regulators better design policies for the industry.

I. INTRODUCTION

The network services and content of the Internet have been evolving since its inception decades ago. Although in the early days, the Internet mostly served elastic traffic and applications, e.g., emails and web page downloading, we have seen significant rise in inelastic traffic, e.g., video and interactive web traffic in recent years. According to [10], from 2007 to 2009, web content traffic had increased from 41.68% to 52%, reaching more than half of the total Internet traffic. Also, when the video streaming giant Netflix moved online a few years ago, its traffic surged immediately. Now, it accounts for up to 32.7% of peak U.S. downstream traffic [2] and its traffic volume is higher than that of BitTorrent [3] applications.

The changes in the Internet content also drive the evolution of the network service providers in many ways. First, besides the traditional single *best-effort* service, i.e., the IP transit service, the new media and its quality of service requirements incentivize the emergence of new SPs and services, e.g., Content Distribution Networks (CDNs) and high-quality video streaming providers, in

the ecosystem. Second, the business relationships between the SPs and CPs and between the SPs themselves have been changing drastically. Because CDNs can reduce the traffic volume from upstream, saving transit costs from their providers, ISPs very often do not charge the CDNs for putting servers in their networks. However, since CDNs can deliver CPs' content faster and more efficiently, they also compete with ISPs for their CP customers. Thus, many ISPs also start to provide their own CDN services besides the IP transit service and charge CDNs via private settlements, e.g., Comcast charges Akamai via a private peering relationship. Among the SPs, ISP settlements were often done bilaterally under either a (zero-dollar) peering or in the form of a customer-provider relationship. Tier-1 ISPs, e.g., Level3, often charge lower tier ISPs for transit services and connect with each other under settlement-free peering. However, peering disputes happened, e.g., the de-peering between Cogent and Level3 in 2005, where the lower tier ISPs that are closer to the content or user side refused to pay for the transit charge. This leads to the recent debate of Network Neutrality [19], which reflects the ISPs' willingness to provide value-added and differentiated services and potentially charge CPs based on different levels of service quality.

To understand why (and how) these changes of the CPs and SPs come about, it is important to first understand the CPs' choices of SPs as well as the market share of the SPs under competition. The goal of this paper is to *model the preferences of the CPs based on their characteristics, e.g., profitability and sensitivity to quality, so as to derive the market share of the SPs with various qualities and prices*. This is of great interests of the SPs, especially when the prices of IP transit [16] and CDN [4] are dropping. Our model will help the SPs to understand the market and make appropriate business decisions, e.g., pricing and quality control, so

as to succeed in the competitive Internet ecosystem. Furthermore, it will also help the regulatory authorities design better policy frameworks for the Internet industry so as to achieve more innovation and competition.

II. THE CONTENT-SERVICE MODEL

We start with a model of the Internet ecosystem that consists a set of Content Providers (CPs) and Service Providers (SPs). The SPs differ by the service qualities they provide and the prices they charge. We model and analyze the CPs' choice of SP based on their own characteristics: how profitable the CP is and how sensitive the CP traffic is to the obtained level of service quality. In essence, this macroscopic model can help us to understand the decision process of these players in the Internet ecosystem and how these decisions may influence their business relationships.

A. The Content and Service Providers

We consider a geographic region with an Internet user base, say L end-users. We denote $(\mathcal{M}, \mathcal{N})$ as a macroscopic model of the Internet ecosystem, consisting of a set \mathcal{M} of SPs and a set \mathcal{N} of CPs. The CPs provide the content that the end-users request, while the SPs provide the network infrastructure for delivering the content to the end-users.

More precisely, our notion of an SP is based on the CPs' point of view. In other words, the services provided by the SPs are for the CPs to reach their customers/users. The definition of an SP is *broader* than the commonly used term: Internet Service Providers (ISPs). ISPs, depending on different taxonomies [8], [7], [12], include 1) eyeball/access ISPs that serve the last-mile for end-users, 2) backbone/Tier 1 ISPs that provide transit services for lower tier ISPs, and 3) content ISPs that serve CPs and host content servers. An SP can be an ISP of any type. Although access and transit ISPs traditionally do not have business relationships with CPs explicitly, with the emergence of video streaming CPs, e.g., Netflix, we have seen more and more CPs' direct or indirect contracts with the access and transit ISPs. For example, Level3 contracted with Netflix for content delivery and Comcast managed to charge Level3 and Akamai via paid-peering contracts (for delivering Netflix's traffic to Comcast's customer base faster) [16]. Although "whether ISPs should be allowed to differentiate services/charges for CPs" is hotly debated under the *network neutrality* [19] argument, legitimate service differentiations will also induce more extensive business relationships among the CPs and ISPs. In general, an SP can be any facilitator

that delivers content to end-users. An important example of an SP that does not even own network infrastructures in the current Internet ecosystem is Akamai [1], which represents the CDNs.

For each SP $I \in \mathcal{M}$, we denote the pair (p_I, q_I) as the SP I 's type. p_I denotes the per unit traffic charge for the CPs to use SP I . q_I denotes the service quality of SP I , e.g., queuing delay or packet loss probability. Without loss of generality, we assume that $q_I \geq 0$ and smaller values of q_I indicate better quality of services. As SPs differ only by price p_I and quality q_I from the CPs' perspective, we can conceptually aggregate the SPs that have the same value pair (p_I, q_I) into a single SP. Similarly, if an SP performs service differentiations, we conceptually treat it as multiple SPs, each with a service class (p_I, q_I) . In fact, our abstraction of an SP I models a market segment that provides a quality level q_I .

To characterize the CPs, we denote u_i as the utility function of CP $i \in \mathcal{N}$. In particular, we define $u_i(p_I, q_I)$ as CP i 's utility when it uses SP I , which depends on the service quality q_I and the per unit traffic charge p_I .

Assumption 1: Any CP i 's utility function $u_i(\cdot, \cdot)$ is non-increasing in both arguments.

Assumption 2: For any set \mathcal{M} of SPs, each CP $i \in \mathcal{N}$ chooses to use an SP, denoted as $I_i \in \mathcal{M}$, that satisfies

$$u_i(p_{I_i}, q_{I_i}) \geq u_i(p_I, q_I), \quad \forall I \in \mathcal{M}.$$

The above assumes that each CP is rational and chooses an SP that provides the highest utility. Technically, there might exist multiple SPs that provide the same amount of utility for the CP. We assume that every CP has certain preference to break the tie and choose one of the SPs. We further denote $\mathcal{N}_I \subseteq \mathcal{N}$ as the set of CPs that choose to use SP I , or the market share of SP I , defined as $\mathcal{N}_I = \{i \in \mathcal{N} : I_i = I\}$.

Lemma 1: Under Assumption 1, $\mathcal{N}_J = \emptyset$, if there exists an SP $I \in \mathcal{M}$ with $q_I < q_J$ and $p_I < p_J$.

Proof: For any $i \in \mathcal{N}$, we have $u_i(p_I, q_I) = \alpha_i(v_i - p_I)e^{-\beta_i q_I} > \alpha_i(v_i - p_J)e^{-\beta_i q_I} > \alpha_i(v_i - p_J)e^{-\beta_i q_J} = u_i(p_J, q_J)$. Thus, by Assumption 2, AP i would choose TP I over TP J . As a result, no AP will choose to use AP J , making $\mathcal{N}_J = \emptyset$. ■

Lemma 1 simply says that if the CPs are rational, no CPs will choose to use an SP (J) that charges a higher price but provides worse quality than an available SP (I) in the market. It also implies that for each market segment of a fixed quality q_I , there would be a single market price p_I , and that better quality services would be priced at higher market prices.

Lemma 2: Under Assumption 1, if \mathcal{M} and \mathcal{M}' are identical except for one SP I with $p'_I > p_I$, then $\mathcal{N}_I(\mathcal{M}') \subseteq \mathcal{N}_I(\mathcal{M})$.

Proof: Let i be an AP $i \in \mathcal{N}_I(\mathcal{M}')$. By Assumption 2, we have $u_i(p'_I, q'_I) \geq u_i(p'_J, q'_J)$ for all $J \in \mathcal{M}' \setminus \{I\}$. Since $p'_I > p_I$, we have $u_i(p_I, q_I) = u_i(p_I, q'_I) \geq u_i(p'_I, q'_I) \geq u_i(p'_J, q'_J) = u_i(p_J, q_J)$ for all $J \in \mathcal{M} \setminus \{I\}$. This implies that AP i will choose TP I over all the TPs, and therefore, $i \in \mathcal{N}_I(\mathcal{M})$. This concludes $\mathcal{N}_I(\mathcal{M}') \subseteq \mathcal{N}_I(\mathcal{M})$. ■

Lemma 2 implies that when an SP unilaterally increases (decreases) its price, fewer (more) CPs will choose to use it. Intuitively, when SP I increases p_I , the utility of each CP in \mathcal{N}_I does not increase. It is possible that some of them move to other SPs which now provide higher utility than SP I . However, no CPs that originally chose other SPs will move to SP I .

Definition 1 (Convexity): The pricing of \mathcal{M} is *convex* if for any SPs $I, J, K \in \mathcal{M}$ with $q_I < q_K < q_J$,

$$p_K \leq \eta p_I + (1 - \eta) p_J,$$

where $\eta = (q_J - q_K)/(q_J - q_I)$.

The above definition is a discrete version of a continuous convex pricing function. Convex pricing often reflects the underlying convex cost where the marginal cost monotonically increases with the level of quality.

Definition 2 (Quasi-Concavity): The utility function u_i is *quasi-concave* if the upper contour sets $\{(p_i, q_i) \in \mathbb{R}_+^2 : u_i(p_i, q_i) \geq u\}$ are convex for all $u \in \mathbb{R}$.

The quasi-concavity of the utility function implies that if two choices (p_1, q_1) and (p_2, q_2) provide at least u amount of utility for CP i , then any linear combination of the choices will induce at least that amount of utility for CP i . In practice, a CP often prefers better quality services until a certain level at which the price becomes a concern. Combined with a convex pricing, a quasi-concave utility function implies this kind of single-peak preference of the CP as follows.

Lemma 3 (Single-Peak Preference): When the pricing of \mathcal{M} is convex and u_i is quasi-concave, for any SPs $I, J \in \mathcal{M}$ with $u_i(p_I, q_I) > u_i(p_J, q_J)$, then $u_i(p_J, q_J) \geq u_i(p_K, q_K)$ if $q_I < q_J \leq q_K$ or $q_I > q_J \geq q_K$.

Proof: We first consider the case $q_I < q_J < q_K$. By Definition 1, we know that $p_J \leq \eta p_I + (1 - \eta) p_K$, where $\eta = (q_K - q_J)/(q_K - q_I)$. By Assumption 1, we have $u_i(p_J, q_J) \geq u_i(\eta p_I + (1 - \eta) p_K, q_J) = u_i(\eta p_I + (1 - \eta) p_K, \eta q_I + (1 - \eta) q_K)$. By Definition 2, we have $u_i(p_J, q_J) \geq u_i(\eta p_I + (1 - \eta) p_K, \eta q_I +$

$(1 - \eta) q_K) \geq \min(u_i(p_I, q_I), u_i(p_K, q_K))$. Since $u_i(p_I, q_I) > u_i(p_J, q_J)$, we must have $u_i(p_J, q_J) \geq u_i(p_K, q_K)$. The derivation of the case $q_I > q_J > q_K$ is similar. ■

Lemma 3 gives a condition under which if a CP prefers a higher (lower) quality SP I over a lower (higher) quality SP J , then it prefers I over any SP whose quality is inferior (superior) to that of J . This condition will help us to understand the collective choice of CPs of different types in the next section.

B. Throughput and Type of the CPs

Although the utility function u_i can be used to model all the characteristics of CP i , the setting does not yet capture the traffic dynamics and the profitability of the CPs. We model CP i 's profitability by denoting v_i as its per unit traffic revenue. This revenue is related to the CP's core business, e.g., online advertising or e-commerce, and we do not assume how it is generated. We denote λ_i as CP i 's throughput function, where $\lambda_i(q_I)$ defines the aggregate throughput of CP i toward its consumers under a quality level q_I . Thus, we model any CP i 's utility as its total profit (profit margin multiplied by the total throughput rate), defined by

$$u_i(p_I, q_I) = (v_i - p_I) \lambda_i(q_I). \quad (1)$$

Assumption 3: For any CP $i \in \mathcal{N}$, $\lambda_i(\cdot)$ is a non-increasing function with an upper-bound $\alpha_i = \lim_{q_i \rightarrow 0} \lambda_i(q_i)$ and lower-bound $\lim_{q_i \rightarrow \infty} \lambda_i(q_i) = 0$.

Assumption 3 says that the throughput will not decrease if a CP uses a better service. λ_i reaches a maximum value of α_i when it receives the best quality $q_i = 0$; while λ_i drops down to zero if the quality deteriorates infinitely, i.e., $q_i = \infty$. If ξ_i percent of the L users would ever be interested in CP i 's content, α_i can be expressed as $\alpha_i = \xi_i L \hat{\theta}_i$, where $\hat{\theta}_i$ denotes the maximum throughput per user under the best service quality. In general, λ_i can be decomposed at a per-user level as a form of $\lambda_i(q_i) = \xi_i L d_i(q_i) \theta_i(q_i)$ [13], where d_i denotes the percentage of active users, θ_i denotes the average throughput per active user. In particular, we consider the following canonical form of the throughput function:

$$\lambda_i(q_I) = \alpha_i e^{-\beta_i q_I}, \quad (2)$$

where CP i 's throughput is characterized by a parameter β_i that captures its sensitivity to the received quality q_I .

Figure 1 illustrates the throughput of two different CPs with parameters $(\alpha_1, \beta_1) = (10, 1.0)$ and $(\alpha_2, \beta_2) =$




Fig. 1. Throughput of different type of CPs.

(6, 0.1) under varying quality levels along the x-axis. Here, we interpret the service quality as queueing delay. CP 1 represents Netflix-type of content that is more sensitive to delay and has a higher maximum throughput rate α_i ; however, CP 2 represents Google-type of content that is less sensitive to delay. We observe that when delay increases from zero, the throughput of delay-sensitive content decreases sharply, while the delay-insensitive content decreases only mildly.

Notice that CPs with the same (β_i, v_i) value pairs would make the same choice of SP. Thus, we can conceptually aggregate them as a single CP. Similar to an SP I representing a market segment, each CP i can be interpreted as a group of CPs with the same characteristics and α_i represents the aggregate maximum traffic intensity, which depends on the number of CPs in the group and the individual traffic intensities. In summary, based on our throughput model, we define

$$u_i(p_I, q_I) = (v_i - p_I)\lambda_i(q_I) = \alpha_i(v_i - p_I)e^{-\beta_i q_I}. \quad (3)$$

Similar to the SPs, we can characterize any CP i as its type, defined by a triple (α_i, β_i, v_i) .

III. CPs' CHOICE AND SPs' MARKET SHARES

When facing a set \mathcal{M} of SPs, each CP i 's best choice I_i depends on the price-quality pairs $\{(p_I, q_I) : I \in \mathcal{M}\}$ and its own characteristics (β_i, v_i) . Notice that since α_i is a linear scaling factor of the throughput, it does not affect the CP's preference among different SPs.

Given any CP i with (β_i, v_i) and a real value u , we define the set $\{(p_I, q_I) : u_i(p_I, q_I) = u\}$ as the indifference set of SPs that provide u amount of utility for CP i . We denote U as the normalized utility defined by $U = u/\alpha_i$ and plot the indifference sets of CP i with $(v_i, \beta_i) = (1.0, 0.5)$ in Figure 2. We vary p_I and q_I on the y-axis and the x-axis. Each point (p_I, q_I) on the




Fig. 2. Indifference sets of CP i , $(v_i, \beta_i) = (1.0, 0.5)$.

plane represents a type of SP. We observe that in order to achieve higher utility, the CP needs a point (p_I, q_I) closer to the origin, which means either the service quality is better, or the charge is cheaper, or both.




Fig. 3. Indifference set for $U = 0.1$ of different CPs.

In Figure 3, we fixed the normalized utility $U = 0.1$ and show the indifference set of different types of CPs. We observe that when β_i increases, which implies that throughput rate becomes more sensitive to quality, the indifference set shifts from right to left, showing that the CPs require a better service quality to keep its utility. Similarly, when v_i decreases, which implies that the profitability weakens, the indifference set shifts from top to bottom, showing that the CPs require a lower pricing by SPs in order to keep its utility.

With this framework, we can analyze and understand the choices made by CPs when there are multiple SPs. To illustrate, we consider the collective choices of the CPs under a market of four SPs. In Figure 4, We fix the qualities to be $(q_1, q_2, q_3, q_4) = (1, 3, 5, 7)$ and the prices to be $(p_1, p_3, p_4) = (0.7, 0.25, 0.1)$ and vary p_2 from 0.3 to 0.6 in the four subfigures from left to right. In each subfigure, we vary β_i on the x-axis




Fig. 4. Shift of market share for four SPs under $(q_1, q_2, q_3, q_4) = (1, 3, 5, 7)$ and $(p_1, p_3, p_4) = (0.7, 0.25, 0.1)$.




Fig. 5. Shift of market share for four SPs under $(q_1, q_2, q_3, q_4) = \kappa(1, 3, 5, 7)$ and $(p_1, p_2, p_3, p_4) = (0.7, 0.4, 0.25, 0.1)$.

and v_i on the y -axis. Each point (β_i, v_i) on the plane represents a type of CP. The CPs located on the top are more profitable and the CPs located on the right are more sensitive to the quality of service. Notice from Figure 1 that a Netflix-type CP i , i.e., $\beta_i = 1$, would obtain around 40% and 5% of its maximum throughput under quality q_1 and q_2 ; however, under q_3 and q_4 , its obtainable throughput almost reaches zero. Thus, CPs with higher value of β_i will more likely choose higher quality SPs. The sets $\mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3$ and \mathcal{N}_4 are shown in yellow, red, green and blue respectively. For example, \mathcal{N}_1 (\mathcal{N}_4) represents the set of CPs that eventually choose the SP that provides the highest (lowest) quality with the highest (lowest) price. For any $I, J \in \mathcal{M}$, we define $\mathcal{N}_{IJ} = \{(\beta_i, v_i) : u_i(p_I, q_I) = u_i(p_J, q_J)\}$ to be the set of CPs that obtain equal utility from I and J . In each sub-figure, we plot \mathcal{N}_{12} and \mathcal{N}_{23} in solid lines and \mathcal{N}_{13} and \mathcal{N}_{24} in dashed lines. Thus, Figure 4 illustrates the *shift of market shares* for these four SPs when we vary the price p_2 of SP 2.

We make the following observations. First, with the in-

crease (decrease) of p_2 , \mathcal{N}_2 decreases (increases) monotonically (by Lemma 2). Second, if we keep increasing (decreasing) p_2 to p_1 (p_3), \mathcal{N}_2 (\mathcal{N}_3) will become empty (by Lemma 1). Third, the upper-right CPs always choose SPs with better qualities. Finally, when $p_2 = 0.4$ or 0.5 , each set \mathcal{N}_i forms a distinct band; however, when $p_2 = 0.3$ and 0.6 , we find \mathcal{N}_3 and \mathcal{N}_2 to be isolated regions respectively. This can be explained by the nature of the pricing of \mathcal{M} and the quasiconcavity of the utility function as follows.

Lemma 4: The utility function $u_i(p_I, q_I) = (v_i - p_I)\lambda_i(q_I)$ is quasiconcave if $\lambda_i(q_I) = \alpha_i e^{-\beta_i q_I}$.

Notice that when $p_2 = 0.4$, the pricing of \mathcal{M} becomes convex and by Lemma 3 and 4, each CP has a single-peak preference among the SPs, where the bands show the preference peaks of the CPs. When $p = 0.3$ or $p = 0.6$, the non-convexity in pricing induces non-single-peak preferences of some CPs. For example, when $p_2 = 0.3$ ($p_2 = 0.6$), we can identify CPs that prefer SP 2 and SP 4 (SP 1 and SP 3) over SP 3 (SP 2), where \mathcal{N}_3 (\mathcal{N}_2) shrinks to be an isolated region.

Let us illustrate the shift of market share when SPs vary their capacity. In Figure 5, we fix the prices $(p_1, p_2, p_3, p_4) = (0.7, 0.4, 0.25, 0.1)$ and qualities $(q_1, q_2, q_3, q_4) = \kappa(1, 3, 5, 7)$, and scale the capacities by $\kappa = 0.2, 0.5, 2$ and 5 from left to right. We observe that when the qualities degrade, CPs' choices move to better quality SPs gradually.

In summary, we presented a framework to help us to analyze (and understand) the CPs' decision on choosing SPs based on each SP I 's quality and price (q_I, p_I) , and the CP i 's profitability and sensitivity to quality (v_i, β_i) .

A. Potential usage for Designing Regulatory Policy

From a regulatory perspective, our framework can also help design desirable policies for the Internet industry. For example, the core debate has centered around the argument whether ISPs should be allowed to provide service differentiation and/or user discrimination, with the notion of user being either the CPs or consumers. Proponents of network neutrality, mostly the CPs, have argued that the Internet has been neutral since its inception and that has been a critical factor in the innovation and rapid growth that has happened on it. Opponents of network neutrality, mostly the ISPs, claim that without some sort of service differentiation, ISPs will lose the incentive to invest in the networks and the end user experience will suffer. Both camps implicitly or explicitly claim that their approach is beneficial for consumers. To analyze the pros and cons of network neutrality, we need to analytically quantify the consumer welfare under different policies. Based on our framework, we can understand the ISPs' strategy to compete for the CPs so as to decide the market shares of the ISPs. After that, we can further determine the performance of the CPs and therefore, the user utility derived from using those CPs. Under various policies, ISPs will be allowed to carry out different levels of service differentiations/qualities, which result different ISP market shares, CPs' choices of ISPs and the corresponding user utilities. By comparing the utilities of all parties in the ecosystem, the policy makers will get a better understanding of the tradeoffs of different policy frameworks for different parties.

IV. RELATED WORK AND FUTURE WORK

Many empirical studies have been tracking the evolution of the Internet using measurements and public data sets [10], [7], [9], [17], [5]. Labovitz et al. [10] measured the inter-domain traffic between 2007 and 2009, and observed the changes in traffic patterns as well as the consolidation and disintermediation of the

Internet core. Gill et al. [9] collected and analyzed traceroute measurements and showed that large CPs are deploying their own wide-area networks. Dhamdhare et al. [7] confirmed the consolidation of the core of the Internet, that brings the content closer to users. Akella et al. [5] used measurements to identify and characterize non-access bottleneck links in terms of their location, latency and available capacity. At the edge of the Internet, Sundaresan et al. [17] studied the network access link performance measured directly from home gateway devices. Our approach to understand the evolution of the Internet ecosystem is based on a theoretic modeling of the CPs choices of SPs.

Many work [6], [8], [12], [15], [18], [11] also focused on the modeling perspective of the Internet evolution. Chang et al. [6] presents an evolutionary model for the AS topologies. Lodhi et al. [11] used an agent-based model to study the network formation of the Internet. Motiwala et al. [15] used a cost model to study the Internet traffic. Valancius et al. combined models and data to study the pricing [18] structure of IP transit market. Faratin et al. [8] and Ma et al. [12] studied the evolution of the ISP settlements. Our approach is more similar to the classic preference theory [14] from microeconomics theory.

Several future directions of this work are as follows:

- 1) The service capacity of the SPs should play a role in the decision process. In particular, a fixed capacity of the SP can only serve a fixed amount of CPs given under certain quality-level guarantee. An interesting research question is how the SPs choose the quality-level to maximize their profits.
- 2) SPs' price will affect the CPs' choice, and therefore, the market share of the SPs. The next step is to study the pricing decisions of the SPs and the resulting market prices of different services.
- 3) To make the model more useful, a future study should also incorporate industrial data to see how the model fits historical prices of the Internet services and project future evolutionary trends of the Internet ecosystem.

V. CONCLUSIONS

In this paper, we model the business decisions and the preferences of the CPs over different SPs based on 1) CPs' profitability v_i , 2) CPs' sensitivity to service quality β_i , 3) SPs' quality q_I and 4) SPs' price p_I . Based on the collective choices of the CPs, we further derive the market share of the SPs for the Internet ecosystem. Our model provides better understanding of the competition

among the SPs, and further guides the SPs' business decisions in terms of pricing and quality control in the complicated and evolving ecosystem.

REFERENCES

- [1] Akamai. <http://www.akamai.com/>.
- [2] Netflix takes up 32.7% of Internet bandwidth, CNN New. <http://edition.cnn.com/2011/10/27/tech/web/netflix-internet-bandwidth-mashable>.
- [3] Netflix Traffic Now Bigger Than BitTorrent. Has Hollywood Won? GIGAOM New. <http://gigaom.com/broadband/netflix-p2p-traffic/>.
- [4] Video Delivery Pricing for Q4 2011. <http://www.cdnpricing.com/>.
- [5] A. Akella, S. Seshan, and A. Shaikh. An empirical evaluation of wide-area Internet bottlenecks. In *Proceedings of the ACM conference on Internet measurement (IMC)*, 2003.
- [6] H. Chang, S. Jamin, and W. Willinger. To peer or not to peer: Modeling the evolution of the Internet's AS-level topology. In *Proceedings of IEEE Infocom*, Barcelona, Spain, 2006.
- [7] A. Dhamdhere and C. Dovrolis. Ten years in the evolution of the Internet ecosystem. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement (IMC 08)*, pages 183–196, Vouliagmeni, Greece, October 2008.
- [8] 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.
- [9] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening Internet topology: natural evolution, unsightly barnacles or contrived collapse? In *Proceedings of the 9th international conference on Passive and active network measurement*, 2008.
- [10] C. Labovitz, D. McPherson, S. Iekel-Johnson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. In *Proceedings of the ACM SigComm*, New Delhi, India, 2010.
- [11] A. Lodhi, A. Dhamdhere, and C. Dovrolis. GENESIS: An agent-based model of interdomain network formation, traffic flow and economics. In *Proceedings of IEEE Infocom*, Miami FL, March 2012.
- [12] 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. *IEEE/ACM Transactions on Networking*, 19(3), June 2011.
- [13] R. T. B. Ma and V. Misra. The Public Option: a non-regulatory alternative to network neutrality. In *Proceedings of 7th ACM Conference on Emerging Network Experiment and Technology (CoNEXT 2011)*, Tokyo, Japan, December 2011.
- [14] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- [15] M. Motiwala, A. Dhamdhere, N. Feamster, and A. Lakhina. Towards a cost model for network traffic. *ACM SIGCOMM Computer Communication Review*, 42(1), January 2012.
- [16] W. Norton. *The Internet Peering Playbook: Connecting to the Core of the Internet*. DrPeering Press, 2011.
- [17] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescap. Broadband Internet performance: A view from the gateway. In *Proceedings of ACM SIGCOMM*, Toronto, Ontario, Canada, August 2011.
- [18] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. Vazirani. How many tiers? pricing in the Internet transit market. In *Proceedings of ACM SIGCOMM*, Toronto, Ontario, Canada, August 2011.
- [19] T. Wu. Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law*, 141, 2005.