

On the Evolution of the Internet Economic Ecosystem

Richard T. B. Ma
ADSC Illinois and National
University of Singapore
tbma@comp.nus.edu.sg

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

Vishal Misra
Columbia University
misra@cs.columbia.edu

ABSTRACT

The evolution of the Internet has manifested itself in many ways: the traffic characteristics, the interconnection topologies and the business relationships among the autonomous components. It is important to understand why (and how) this evolution came about, and how the interplay of these dynamics may affect future evolution and services. We propose a network aware, macroscopic model that captures the characteristics and interactions of the application and network providers, and show how it leads to a market equilibrium of the ecosystem. By analyzing the driving forces and the dynamics of the market equilibrium, we obtain some fundamental understandings of the cause and effect of the Internet evolution, which explain why some historical and recent evolutions have happened. Furthermore, by projecting the likely future evolutions, our model can help application and network providers to make informed business decisions so as to succeed in this competitive ecosystem.

Categories and Subject Descriptors

C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks—*Internet*; C.4 [Performance of Systems]: [Modeling techniques, Performance attributes]

Keywords

Internet Evolution; Economics; IP Transit; Content Delivery

1. INTRODUCTION

The Internet has been and is still changing unexpectedly in many aspects. Started with elastic traffic and applications, e.g., emails and webpage downloading, we have seen significant rise in inelastic traffic, e.g., video and interactive web traffic, across the Internet. According to [19], from 2007 to 2009, web content traffic had increased from 41.68% to 52%, reaching more than half of the total Internet traffic. From a network perspective, the Internet originated from government-owned backbone networks, i.e., the ARPANET, and then evolved to a network of commercial Autonomous Systems (ASes) and Internet Service Providers (ISPs). Meanwhile, ISPs formed a hierarchical structure and were classified by tiers, with higher tier ISPs cover larger geographic regions and provide transit service for smaller/lower tier ISPs. However, recent study [17] has reported that large content providers, e.g., Google and Microsoft, are deploying their own wide-area networks so as to bring content closer to users and bypassing Tier-1 ISPs on many paths. This is known as the *flattening phenomenon* of the Internet topology.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). IW3C2 reserves the right to provide a hyperlink to the author's site if the Material is used in electronic media.
WWW 2013, May 13–17, 2013, Rio de Janeiro, Brazil.
ACM 978-1-4503-2035-1/13/05.

Changes in the content or network topology do not happen independently. Rather, they are driven by the changes in the business relationships among the players in the Internet ecosystem. Not surprisingly, we have observed dramatic changes in the business relationships between the content providers and the ISPs and among the ISPs themselves. Traditionally, 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., Level 3 [5], often charge lower tier ISPs for transit services and connect with each other under settlement-free peering. However, the Tier-1 ISPs do not have any guarantee in their profitability as the Internet evolves. For instance, we have seen exponential decrease (around 20% a year) in IP transit prices [23]. Also, peering disputes happened, e.g., the de-peering between Cogent [3] and Level 3 in 2005, where the lower tier ISPs that are closer to content or users refused to pay for the transit charge. This leads to the recent debate of network neutrality [26], which reflects the ISPs' willingness to provide value-added and differentiated services and potentially charge content providers based on different levels of service quality.

The situation is further complicated by the emergence of new players in the ecosystem: Content Delivery Networks (CDNs), e.g., Akamai [1] and Limelight [6], and high-quality video streaming providers, e.g., Netflix [7]. From content providers' perspective, CDNs can deliver their content faster and more efficiently; from local ISPs' perspective, CDNs can reduce the traffic volume from upstream, saving transit costs from their providers. Very often, ISPs do not charge the CDNs for putting servers in their networks. 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 [8] and its traffic volume is higher than that of BitTorrent [9] applications. Netflix used Limelight, one of the biggest CDNs, for content delivery, and later, the Tier-1 Level 3 also obtained a contract to deliver Netflix's traffic. Since most of the Netflix customers are based in the U.S., they often use Comcast, the biggest access ISP, as the last-mile access provider. Interestingly, Comcast managed to enter a so-called paid-peering relationship [16] with Level 3 and Limelight, under which the Tier-1 ISP and the CDN have to pay the access ISP for higher bandwidth on the last mile connection. This has totally *reversed* the nominal customer-provider relationship where the Tier-1 ISP was the service provider and should have received payment for connectivity.

It is important to understand how these changes come about, and what the driving factors are behind these changes. In this work, we model the Internet evolution from a macroscopic view that captures the cause and effect of the evolution of the individual players in the ecosystem. Our model expends the traditional view of a single best-effort service model to capture multiple value-added services in the Internet. The main approaches and contributions are as follows.

- We model the preferences and business decisions of the application providers for purchasing Internet services, based on the application characteristics and the price and quality of the transport services (Section 2).
- We characterize the market price and the market share of the Internet transport services by using general equilibrium theory in economics (Section 3).
- We analyze the driving forces of the evolution of the Internet economic ecosystem (Section 4), which provide qualitative answers (Section 4.6) to questions like: 1) *Why have the IP transit prices been dropping?* 2) *Why have the CDNs emerged in the ecosystem?* 3) *Why has the pricing power shifted to the access ISPs?* 4) *Why are the large content providers building their own wide-area networks toward users?*
- We incorporate Internet price and capacity data into our model, and quantitatively fit historical prices and project the future evolution of the ecosystem and its price trends (Section 5).
- We demonstrate how our model can help the network providers to make business decisions, e.g., capacity expansion and peering decisions, based on the future price projections under various scenarios (Section 5.3).

Our paper sheds new light on the macroscopic evolution of the Internet economic ecosystem and concretely identifies the driving factors of such an evolution. In particular, our model provides a tool to analyze and project the evolutionary trends of the ecosystem. The fundamental understanding of the preferences of application providers and the market equilibrium of the Internet services will also help the business decisions of the application and network transport providers to succeed in this competitive ecosystem.

2. THE MACROSCOPIC AP-TP MODEL

We start with a macroscopic model of the Internet ecosystem that consists a set of Application Providers (APs) and Transport Providers (TPs). The TPs differ by their service qualities and the prices they charge. We model and analyze the APs' choice of TP based on their own characteristics: how profitable the AP is and how sensitive the AP 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.

2.1 The Application and Transport Providers

We consider an Internet service market of a geographic region and denote $(\mathcal{M}, \mathcal{N})$ as a macroscopic model of the ecosystem, consisting of a set \mathcal{M} of TPs and a set \mathcal{N} of APs. The APs provide the content/service for the Internet end-users; the TPs provide the network infrastructure for delivering the APs' data to their end-users.

Our notion of an AP broadly includes content providers, e.g., Netflix, online services, e-commerce, and even cloud services, e.g., Amazon EC2 [2]. Our notion of a TP is based on the APs' point of view. In other words, the transport services provided by the TPs are for the APs to reach their customers/users. The scope of a TP is *broader* than an ISP, and it includes CDNs. ISPs, depending on different taxonomies [16, 15, 21], 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 APs and host content servers. A TP can be an ISP of any type. Although access and transit ISPs traditionally do not have business relationships with APs explicitly, with the emergence of

video streaming APs, e.g., Netflix, we have seen more and more APs' direct or indirect contracts with the access and transit ISPs. For example, Level 3 contracted with Netflix for content delivery and Comcast managed to charge Level 3 and Limelight via paid-peering contracts (for delivering Netflix's traffic to Comcast's customer base faster) [23]. Although "whether ISPs should be allowed to differentiate services/charges for APs" is hotly debated under the network neutrality [26] argument, legitimate service differentiations will also induce more extensive business relationships among the APs and ISPs. In general, a TP can be any facilitator that delivers content to end-users. An important example of a TP that does not even own network infrastructures in the current Internet ecosystem is Akamai [1], which represents the CDNs.

We characterize each TP $I \in \mathcal{M}$ by its type, denoted as a triple (p_I, q_I, ν_I) . p_I denotes the per unit traffic charge for the APs to use TP I . q_I denotes the service quality of TP I , e.g., queueing 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. ν_I denotes the bandwidth capacity of TP I . We characterize each AP $i \in \mathcal{N}$ by its utility function $u_i(\cdot)$. In particular, we define $u_i(p_I, q_I)$ as AP i 's utility when it uses TP I , which depends on the service quality q_I and the per unit traffic charge p_I .

Assumption 1. $u_i(\cdot, \cdot)$ is non-increasing in both arguments.

Assumption 2. For any set \mathcal{M} of TPs, each AP $i \in \mathcal{N}$ chooses to use a TP, 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 AP is rational and chooses a TP that provides the highest utility. Technically, there might exist multiple TPs that provide the same amount of utility for the AP. We assume that every AP has certain preference to break the tie and choose one of the TPs. We further denote $\mathcal{N}_I \subseteq \mathcal{N}$ as the set of APs that choose to use TP I , or the market share of TP I , defined as $\mathcal{N}_I = \{i \in \mathcal{N} : I_i = I\}$. Based on Assumption 1 and 2, if two TPs I, J have the same quality, i.e., $q_I = q_J$, then they have to price equally, i.e., $p_I = p_J$; otherwise, the one with higher price will not obtain any market share. As TPs differ only by price p_I and quality q_I from the APs' perspective, we aggregate the TPs that have the same value pair (p_I, q_I) into a single TP with a capacity that equals the summation of individual TPs' capacity. Similarly, if a TP performs service differentiations, we conceptually treat it as multiple TPs, each with a service class (p_I, q_I) and the corresponding capacity ν_I . More precisely, our abstraction of a TP I models a competitive market segment that provides a quality level q_I and has a total capacity ν_I .

2.2 Throughput and Types of the APs

Although the utility function u_i can be used to model all the characteristics of AP i , the setting does not yet capture the traffic dynamics and the profitability of the APs. We model AP i 's profitability by denoting v_i as its per unit traffic revenue. This revenue is related to the AP's core business, e.g., online advertising or e-commerce, and we do not assume how it is generated. We denote $\lambda_i(\cdot)$ as AP i 's throughput function, where $\lambda_i(q_I)$ defines the aggregate throughput of AP i toward its consumers under a quality level q_I . Thus, we model any AP 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 AP $i \in \mathcal{N}$, $\lambda_i(\cdot)$ is a non-increasing function with $\alpha_i = \lim_{q_i \rightarrow 0} \lambda_i(q_i)$ and $\lim_{q_i \rightarrow \infty} \lambda_i(q_i) = 0$.

Assumption 3 says that the throughput will not decrease if an AP uses a better service. λ_i reaches a maximum value of α_i when it receives the best quality $q_i = 0$ and decreases to zero if the quality deteriorates infinitely, i.e., q_i tends to $+\infty$. 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 AP i 's throughput is characterized by a parameter β_i that captures its sensitivity to the received quality q_I .




Figure 1: Throughput of different type of APs.

Figure 1 illustrates the throughput of two APs with parameters $(\alpha_1, \beta_1) = (10, 1.0)$ and $(\alpha_2, \beta_2) = (6, 0.1)$ under varying service qualities, interpreted as network delays in this case, along the x-axis. AP 1 represents a Netflix-type of application that is more sensitive to delay and has a high maximum rate $\alpha_1 = 10$ Mbps; however, AP 2 represents a Google-type of query application that is less sensitive to delay. We observe that when delay increases, the throughput of delay-sensitive application decreases sharply, while the delay-insensitive application decreases only mildly.

Because α_i is just a linear scaling factor of the throughput, it does not affect the AP's preference over different TPs. Consequently, APs with the same (β_i, v_i) value pairs will choose the same TP; and therefore, we can conceptually aggregate them as a single AP. Similar to a TP I representing a market segment, each AP i can be interpreted as a group of APs with the same characteristics and α_i represents the aggregate maximum traffic intensity, which depends on the number of APs in the group and the individual traffic intensities. Although α_i does not play a role in the AP's decision of choosing TPs, we will see later that α_i reflects the demand of the APs and affects the market prices of the TPs. 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 each TP I 's type (p_I, q_I, ν_I) , we can characterize any AP i 's type as another triple (α_i, β_i, v_i) .

2.3 APs' Choice of Transport Providers

When facing a set \mathcal{M} of TPs, each AP 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) . The APs' choices satisfy the following results.

Theorem 1. For a fixed set \mathcal{M} and any two APs i and j with $\beta_j \geq \beta_i$ and $v_j \geq v_i$, their chosen service qualities satisfy $q_{I_i} \geq q_{I_j}$.

Theorem 1 says that if an AP j is more profitable and more sensitive to service quality than another AP i , then the chosen quality of AP j will be at least as good as that of AP i . This property holds regardless how the services are priced.

Theorem 2. For any $\kappa_1, \kappa_2, \kappa_3 > 0$, and system $(\mathcal{M}, \mathcal{N})$, we define a scaled system $(\mathcal{M}', \mathcal{N}')$ as $\mathcal{M}' = \{(\kappa_1 p_I + \kappa_2, q_I / \kappa_3, \nu_I) :$

$I \in \mathcal{M}\}$ and $\mathcal{N}' = \{(\alpha_i, \kappa_3 \beta_i, \kappa_1 v_i + \kappa_2) : i \in \mathcal{N}\}$, then system $(\mathcal{M}', \mathcal{N}')$ satisfies $\mathcal{N}'_I(\mathcal{M}', \mathcal{N}') = \mathcal{N}_I(\mathcal{M}, \mathcal{N})$ for all $I \in \mathcal{M}$.

Theorem 2 says that if 1) the AP profitability v_i and the TP price p_I are linearly scaled in the same way, and/or 2) the quality q_I of the TPs and the sensitivity β_i of the APs scale inversely at the same rate, then the APs' choices of TP will not change. This result will help us normalize different systems and make a fair comparison of various solutions.

Theorem 3. For any $\kappa > 0$ and a fixed set \mathcal{N} of APs, let $\mathcal{M}' = \{(p_I, \kappa q_I, \nu_I) : I \in \mathcal{M}\}$, then for all $i \in \mathcal{N}$, 1) $q_{I'_i} \leq \kappa q_{I_i}$ if $\kappa > 1$ and 2) $q_{I'_i} \geq \kappa q_{I_i}$ if $\kappa < 1$.

Theorem 3 says that if all the qualities in the market deteriorate ($\kappa > 1$) linearly at the same rate, APs will not use worse quality TPs than before. The opposite is also true: when qualities improve linearly, APs will not use better quality TPs than before.

3. MARKET EQUILIBRIUM

In this section, we start with the definition of a market equilibrium, under which the prices of the TPs are stable and the claimed service qualities can be achieved when APs choose their best TPs. We then proceed to characterize the market equilibrium and calculate the equilibrium prices.

3.1 The Existence of Market Equilibrium

Although any TP I claims to provide service quality q_I , it cannot keep its promise if more APs choose this TP than its capacity can support. We model the achieved quality $Q_I(\lambda_I, \nu_I)$ as a function of the actual throughput λ_I going through I and its capacity ν_I .

Assumption 4. The achieved quality $Q_I(\lambda_I, \nu_I)$ for any TP $I \in \mathcal{M}$ is non-decreasing in λ_I and non-increasing in ν_I .

Definition 1. A set $\mathcal{X} \subseteq \mathcal{N}$ of APs is feasible for TP I with quality q_I , if $Q_I(\lambda_I(\mathcal{X}), \nu_I) \leq q_I$, where $\lambda_I(\mathcal{X}) = \sum_{i \in \mathcal{X}} \lambda_i(q_I)$ defines the induced throughput of the set \mathcal{X} of APs under quality q_I .

In a market \mathcal{M} of TPs, each TP would adjust its strategies to accommodate its customer APs' traffic demand and keep its service quality promise. For example, if the current capacity of TP I cannot support quality q_I , it might 1) expend its capacity ν_I , 2) increase price p_I , or 3) reduce the quality level q_I . Next, we define a market equilibrium where the APs' demand are feasible and the TPs' strategies are stable.

Definition 2. Let p_I^{min} be the cost (or minimum price) of TP I . Let \mathcal{M}' be identical to \mathcal{M} except for $p'_I \neq p_I$ for some $I \in \mathcal{M}$ and \mathcal{N}'_I be the set of APs choosing TP I under \mathcal{M}' . A system $(\mathcal{M}, \mathcal{N})$ forms a market equilibrium if 1) all APs' aggregate demands are feasible, i.e., $Q_I(\lambda_I(\mathcal{N}_I), \nu_I) \leq q_I, \forall I \in \mathcal{M}$, and 2) each price p_I maximizes the utilization of capacity for acceptable throughput at TP I , i.e., for any $p'_I \geq p_I^{min}$ with the corresponding \mathcal{N}'_I satisfying $Q_I(\lambda_I(\mathcal{N}'_I), \nu_I) \leq q_I, \lambda_I(\mathcal{N}'_I) \leq \lambda_I(\mathcal{N}_I)$.

One way to understand the above definition of a market equilibrium is that given a set \mathcal{N} of APs and a set $\{q_I : I \in \mathcal{M}\}$ of service qualities for them to choose from, the price p_I and capacity ν_I of each market segment should be consistent in that 1) when the APs make their choices of TP, their expected service quality can be achieved and, 2) the capacities of the TPs are not under-utilized, unless the charge p_I reaches the TP's cost p_I^{min} . If APs' quality expectations are not fulfilled, their choices of TP will change. Furthermore, if capacity ν_I is under-utilized with $p_I > p_I^{min}$, then the market segment I is not correctly priced. That being said, we assume that none of the market segment is controlled by a monopoly,

which might want to under-utilize capacity and keep a higher price for profit-maximization. We will summarize and discuss the limitations of our model in Section 5.4. The interesting aspect here is that although p_I , like all other prices, mainly depends on the supply ν_I and the demand \mathcal{N}_I of the APs, all the TPs (or market segments) are correlated, which serve substitutions for the APs.

In practice, the TPs might not have enough capacities to accommodate all APs. As a result, market prices will rise and some APs cannot afford the prices and will not use any of the TPs. However, under Assumption 2, each AP needs to choose a TP even it cannot afford to use any of the TPs, so a market equilibrium might not exist under this assumption. To fix this minor technical issue, we make the following assumption to allow any AP not to use any of the TPs if they all induce negative utilities.

Assumption 5. *There always exists a dummy TP $D \in \mathcal{M}$ with quality $q_D = \infty$ and price $p_D = 0$.*

By Assumption 3, quality q_D always induces zero throughput for any AP, and therefore, the dummy TP guarantees a zero utility and can accommodate as many APs as possible in equilibrium. Effectively, the set \mathcal{N}_D models the APs that cannot afford to use any TP in the market in reality.

Theorem 4. *For any fixed set \mathcal{N} of APs and any set \mathcal{M} of TPs with fixed values of p_I^{min} , q_I and ν_I for all $I \in \mathcal{M}$, there exists a set $\{p_I : I \in \mathcal{M}\}$ of prices that makes $(\mathcal{M}, \mathcal{N})$ a market equilibrium.*

Although TPs might be able to adopt new technologies to improve or differentiate their services, the quality that they can provide is often physically constrained by the nature of the TP, for example, if a TP is a Tier 1 ISP, it cannot guarantee end-to-end delays for the customers unless the access ISP's link is not congested. Similarly, although TPs might execute a long-term capacity planning, the supply of capacity does not change in a small time scale. Compared to service quality and capacity, market prices change more frequently and easily. Theorem 4 says that even in a small time-scale where prices adapt to market conditions, prices might still converge to an equilibrium, which reflects the short-term market structure of the Internet ecosystem.

3.2 Characteristics of a Market Equilibrium

In theory, one might find multiple sets of prices that make $(\mathcal{M}, \mathcal{N})$ a market equilibrium. For example, from any existing equilibrium, one might find a TP I such that with only a small change in p_I , no APs will change their choices. This new price also constitutes a market equilibrium. In practice, these price differences can happen by two reasons. First, even without a monopoly in a market segment, oligopolistic providers might implicitly collude on the price so that they keep a relatively high price simultaneously. When one of them starts to reduce price, the price of that segment will converge to a lower price. Second, the preferences of the APs are quite different so that the price change in one segment might not affect the demand choices of the APs.

Definition 3. *A market equilibrium $(\mathcal{M}, \mathcal{N})$ is competitive if there does not exist any $p_I^{min} \leq p'_I < p_I$ with the corresponding \mathcal{N}'_I satisfying $Q_I(\lambda_I(\mathcal{N}'_I), \nu_I) \leq q_I$.*

If the AP types are very diverse or each market segment consists of many competing providers, one can focus on the above definition of a competitive market equilibrium. Technically, a competitive market equilibrium might not exist, since the minimum price might not exist when all the feasible equilibrium prices form an open set. However, prices in practice have a minimum unit, e.g., one cent, and we can always find such a competitive market equilibrium.

Notice that our model is not limited to competitive market equilibria, i.e., if a segment I is not competitive enough, we can use a higher price for p_I . As a result, competitive equilibrium prices might be biased downward if the real market structure is not perfectly competitive; nevertheless, our qualitative results do not depend on whether the market equilibrium is competitive or not.

Theorem 5. *Let $\mathcal{N}' = \{(\kappa\alpha_i, \beta_i, \nu_i) : i \in \mathcal{N}\}$ and $\mathcal{M}' = \{(p_I, q_I, \kappa\nu_I) : I \in \mathcal{M}\}$ for some $\kappa > 0$. If $(\mathcal{M}, \mathcal{N})$ is a market equilibrium and the quality function $Q_I(\cdot, \cdot)$ s are homogenous of degree 0, i.e., $Q_I(\lambda_I, \nu_I) = Q_I(\kappa\lambda_I, \kappa\nu_I)$, $\forall \kappa > 0, I \in \mathcal{M}$, then $(\mathcal{M}', \mathcal{N}')$ is a market equilibrium too.*

Theorem 5 says that if the quality only depends on the ratio of the incoming traffic rate and the capacity, then when the number of APs (and their traffic intensity) and the capacities scale at the same speed, the original market equilibrium prices will remain in equilibrium. If we consider the queueing delay as the quality metric, because of statistical multiplexing, the average queueing delay reduces when both arrival rate and service rate scales up at the same rate. In this case, Theorem 5 also implies that each TP I can accept more and more traffic for a fixed delay q_I , and as a consequence, the market prices will move downward in a new equilibrium.

3.3 Calculating Market Equilibrium Prices

We denote μ_I as the maximum throughput that TP I can accept when it can still fulfill the quality q_I , defined as

$$\mu_I = \arg \max_{\lambda_I} Q_I(\lambda_I, \nu_I) \leq q_I. \quad (4)$$

For instance, if the quality metric is the average queueing delay under M/G/1 systems and TP I implements a FIFO scheduling policy, by the Pollaczek-Khinchine mean formula, $Q_I(\lambda_I, \nu_I) = \frac{\lambda_I}{\nu_I - \lambda_I} E[R]$, where $E[R]$ denotes the expected residual service time of jobs. If we want λ_I to be feasible, we need $\frac{\lambda_I}{\nu_I - \lambda_I} E[R] \leq q_I \Rightarrow \lambda_I \leq \frac{q_I}{E[R] + q_I} \nu_I = \mu_I$. We define $\eta_I = \mu_I / \nu_I$ as the maximum acceptable throughput per unit capacity, or the conversion factor from raw capacity to achievable throughput. Notice that given a fixed capacity ν_I , the smaller delay TP I wants to provide, the smaller maximum amount of traffic it can accept. For the M/G/1 case, η_I tends to 0 when the required quality q_I tends to 0, which shows a convex cost structure for the TP.

Based on the monotonicity of Q_I (Assumption 4), a market equilibrium can be characterized by using μ_I s as follows.

Definition 4. *A system $(\mathcal{M}, \mathcal{N})$ forms a market equilibrium if for all TP I , 1) $\lambda_I(\mathcal{N}_I) \leq \mu_I$, and 2) there does not exist $p'_I \geq p_I^{min}$ with the corresponding \mathcal{N}'_I satisfying $\lambda_I(\mathcal{N}_I) < \lambda_I(\mathcal{N}'_I) \leq \mu_I$.*

Based on the above alternative definition of a market equilibrium, we can calculate the competitive equilibrium prices without evaluating Q_I repeatedly as follows.

Calculate Price Equilibrium $(\mathcal{N}, \{p_I^{min}, q_I, \nu_I : I \in \mathcal{M}\})$

1. Set $p_I = \infty$ for all TP $I \in \mathcal{M}$;
 2. Calculate μ_I for all TP $I \in \mathcal{M}$ based on q_I and Q_I ;
 3. while there exists $p'_I \in [p_I^{min}, p_I)$ such that $\lambda_I(\mathcal{N}_I) \leq \lambda_I(\mathcal{N}'_I) \leq \mu_I$
 4. set $p_I = p'_I$;
 5. return $\{p_I : I \in \mathcal{M}\}$;
-

In the above algorithm, we do not restrict which TP I to choose in step 3 if multiple TPs satisfy the condition. However, any sequence of updates will make the price vector converge, because

each p_I will only be decreasing monotonically until convergence. Similarly, we can also set $p_I = p_I^{min}$ for all TPs, and the price vector will increase monotonically until convergence.

Based on Theorem 2 and 5, we also have the following result.

Corollary 1. Let $\mathcal{N}' = \{(\kappa\alpha_i, \kappa_3\beta_i, \kappa_1v_i + \kappa_2) : i \in \mathcal{N}'\}$ and $\mathcal{M}' = \{(\kappa_1p_I + \kappa_2, q_I/\kappa_3, \nu_I') : I \in \mathcal{M}'\}$ for positive $\kappa, \kappa_1, \kappa_2, \kappa_3$ with $\mu_I' = \kappa\mu_I$ for all $I \in \mathcal{M}'$. If $(\mathcal{M}, \mathcal{N})$ is a market equilibrium, then $(\mathcal{M}', \mathcal{N}')$ is a market equilibrium.

Although the prices of the TPs influence the APs' choices, which further affect the capacity utilization of the TPs, equilibrium prices are the fixed points in which both the APs' choices and the TPs' prices do not change. However, external factors could move the resulting equilibrium. In the next section, we will study these fundamental driving forces for the evolution of the Internet economic ecosystem. By understanding these factors, we will know why the market prices change and why certain evolutions happen.

4. PRICE DYNAMICS IN EQUILIBRIUM

In this section, we look deeper into the qualitative dynamics of the equilibrium market prices. In particular, we explore how the different characteristics of the APs and the TPs can affect the market prices in equilibrium.

4.1 Evaluation Setting

Each AP i is characterized by three parameters (α_i, β_i, v_i) ; each TP I is characterized by three parameters (p_I, q_I, ν_I) . To make a fair comparison between equilibrium prices under different settings, we carefully normalize the system parameters as follows. We define $v_{max} = \max\{v_i : i \in \mathcal{N}'\}$, $\beta_{max} = \max\{\beta_i : i \in \mathcal{N}'\}$, and $p_{min} = \min\{p_I^{min} : I \in \mathcal{M}'\}$. Based on Theorem 2, we normalize any system $(\mathcal{M}, \mathcal{N})$ by factors $\kappa_1 = 1/(v_{max} - p_{min})$, $\kappa_2 = p_{min}/(v_{max} - p_{min})$, and $\kappa_3 = 1/\beta_{max}$. As a result, we normalize each β_i or v_i within the interval $[0, 1]$ and the equilibrium prices will also be scaled accordingly with $[0, 1]$. If p_I^{scaled} is the derived market equilibrium price in the normalized system, we can recover the real market price p_I as

$$p_I = (v_{max} - p_{min})p_I^{scaled} + p_{min}.$$

When the normalized price p_I^{scaled} tends to 0, it reflects that the real market price p_I goes down to the cost p_{min} ; when p_I^{scaled} tends to 1, it reflects that the real market price p_I goes to the maximum AP profitability v_{max} . We describe the TPs' capacity in terms of the maximum acceptable rates μ_I s. We define $\alpha = \sum_{i \in \mathcal{N}} \alpha_i$, $\mu = \sum_{I \in \mathcal{M}} \mu_I$ and the ratio $\rho = \mu/\alpha$. Based on Corollary 1, any price equilibrium sustains when α_i s and μ_I s scale at the same rate. Thus, we normalize the APs' aggregate maximum traffic intensity α to be 1. We define $\sigma_I = \mu_I/\mu$ as the capacity share of TP I , and under the normalized system, each TP I has $\mu_I = \sigma_I\rho$.

After the above normalization, we can describe any system by the following four parameters:

1. a set of qualities $\{q_I : I \in \mathcal{M}'\}$,
2. the normalized aggregate capacity ρ ,
3. the distribution of α_i over the domain $[0, 1]^2$ of (β_i, v_i) ,
4. the capacity distribution $\{\sigma_I : I \in \mathcal{M}'\}$.

We focus on three different quality types: 1) q_A , the highest quality for real-time content delivery, 2) q_B , medium quality, mostly for web applications, and 3) q_C , the best-effort quality, mostly for elastic traffic. As analyzed in [25], IP transit markets will be

quite efficient if two tiers of services are provided; thus, q_B and q_C can be considered as the higher and lower tier services of such an IP transit market. To differentiate the three qualities, we set $q_A : q_B : q_C = 1 : 5 : 25$. We vary ρ from 0 to 1, where the system's total capacity varies from extremely scarce to abundant. We discretize the AP domain with 50 levels of v_i and β_i , which forms 2500 types of APs. We assume that APs' profitability and quality-sensitivity follow probability distributions F_v and F_β respectively, and α_i follows the joint distribution of F_v and F_β . We use the various distributions in Figure 2 for F_v and F_β . For instance, when




Figure 2: Common distributions: geometric, uniform, reversed geometric, binomial with $p = 0.2, 0.5$ and 0.8 .

a geometric distribution Geo is used to describe F_β , it models the scenario where most of the AP traffic are elastic and the amount of quality-sensitive traffic decreases exponentially with its sensitivity level β_i . The binomial distributions $BN(p)$ are often used to approximate a normal distribution of the profitability v_i , or quality sensitivity β_i , where p determines the mean value.

4.2 Impact of TP Capacity on Prices

In this subsection, we study how the capacities of the TPs affect the equilibrium prices. We initially set $(q_A, q_B, q_C) = (0.2, 1, 5)$. We will evaluate how the quality may impact the equilibrium prices in the next subsection.




Figure 3: Shift in market prices as μ_B varies: with $(q_A, q_B, q_C) = (0.2, 1, 5)$, $\mu_A = 0.05$ and $\mu_C = 0.25$.

In Figure 3, we fix $\mu_A = 0.05$, $\mu_C = 0.25$ and vary μ_B from 0 to 1 along the x-axis. The three sub-figures show the equilibrium prices when α_i follows the joint distributions of $(F_\beta, F_v) = (Uni, Uni)$, (Geo, Uni) and $(Geo, BN(0.5))$ respectively. We observe that when μ_B is scarce, equilibrium price p_B is close to (but strictly less than) the price p_A of its upper class TP. When μ_B increases, p_B diverges from p_A and moves to the price p_C of its lower class TP. When μ_B becomes abundant, its market price goes




Figure 4: Shift in market prices as ρ varies: with $(q_A, q_B, q_C) = (0.2, 1, 5)$ and $(F_\beta, F_v) = (Geo, Uni)$.

down to the minimum price after p_C . In general, when the capacity of a particular TP, i.e., μ_B , increases, it drives all equilibrium prices down; however, the prices of higher quality TPs, e.g., p_A , might not go down to the minimum price.

In the rest of this section, we often use $F_\beta = Geo$, which models the case where more APs were elastic, and $F_v = BN(0.5)$, which approximates that the AP profitability follows a normal distribution centered at $v_i = 0.5$. Note that our qualitative results do not depend on these settings.

In Figure 4, we vary the system capacity ρ from 0 to 1 along the x-axis. α_i follows the joint distribution $(F_\beta, F_v) = (Geo, BN(0.5))$. The sub-figures show the equilibrium prices when the capacity ratio $\sigma_A : \sigma_B : \sigma_C$ equals 1 : 3 : 5, 1 : 1 : 1 and 5 : 3 : 1 respectively. In all three cases, when the total capacity ρ is small, all equilibrium prices are very close and high. When we increase ρ , all market prices drop. By comparing the price curves across the three subfigures, we observe that when the capacity share of the higher class TP is smaller (the left subfigure), 1) the three market prices differ more from each other, 2) p_C drops faster, and 3) all the prices drop to the minimum price faster than the other two cases. Because price differences exist in practice, we will use $\sigma_A : \sigma_B : \sigma_C = 1 : 3 : 5$ in the rest of this section.

Lessons (the TP capacity effects on prices) leaned:

- Capacity expansion drives market prices down.
- The capacity expansion of a particular TP I would affect not only its own price p_I , but also other TPs' prices, due to the substitution effect of TP I to other TPs.
- When TP I 's capacity share σ_I is small (big), its market price p_I is close to the price of its next higher (lower) class TP.

4.3 Impact of TP Quality on Prices

Let us explore how the quality q_I of the TPs may affect the equilibrium prices. We use the setting that the capacity distribution follows $\sigma_A : \sigma_B : \sigma_C = 1 : 3 : 5$ and α_i follows the joint distribution $(F_\beta, F_v) = (Geo, BN(0.5))$.

In Figure 5, we keep the quality ratio $q_A : q_B = q_B : q_C = 1 : 5$ and use $(q_A, q_B, q_C) = \kappa(0.2, 1, 5)$, where κ equals 0.2, 1 and 5 in the three subfigures. We vary the system capacity ρ from 0 to 1 along the x-axis. We observe that when all the TPs improve their quality by the same ratio, i.e., $\kappa = 0.2$, the market prices of the TPs are very close; when all the TPs degrade their quality by the same ratio, i.e., $\kappa = 5$, the market prices of the TPs diverge greatly. This observation can be explained by Theorem 3. When κ decreases and all qualities are improved, more APs will choose lower class TPs, which move the prices of the lower class TPs upward and the prices of upper class TPs downward. As a result, all TPs prices will move closer. On the other hand, when κ increases and all qualities are




Figure 5: Shift in market prices as ρ varies: with $(q_A, q_B, q_C) = \kappa(0.2, 1, 5)$ where $\kappa = 0.2, 1$ and 5 .




Figure 6: Shift in market prices as ρ varies: with $q_B = 1$, $q_A : q_B = q_B : q_C = 1 : \kappa$ where $\kappa = 2, 5, 10$.

degraded, more APs will choose to upper class TPs, which move the prices of the upper class TPs upward and the prices of lower class TPs downward. This will further diverge the price differences among the TPs with different qualities.

In Figure 6, we keep $q_B = 1$ and vary the quality ratio $q_A : q_B = q_B : q_C = 1 : \kappa$, where κ equals 2, 5 and 10. We observe that the price differences are positively correlated with the quality ratio. In particular, when quality ratio is high, e.g., $\kappa = 10$, the price of the lowest class TP, i.e., p_C , drops earlier and sharper when the total capacity ρ expands. At the same time, higher class TPs can still maintain a non-zero market price even after p_C drops down the minimum price. The general trend is that when the quality ratio keeps increasing, the price curves will move higher and toward the left. In the rest of this section, we will often use the quality ratio 1 : 5 and $(q_A, q_B, q_C) = (0.2, 1, 5)$ for our evaluations. Again, our qualitative results do not depend on this setting.

Lessons (the TP quality effects on prices) leaned:

- The market prices of the TPs would be close to (far from) one another if the quality ratio is small (big) or/and the overall qualities of the market are high (low).
- In reality, the qualities provided by the TPs are becoming better and better, which implies that market prices for different services might converge.
- High-end market segments can still maintain a price difference if they can differentiate their quality from the lower class TPs substantially.

Next, we will see that the TP price differences also depend on the demand side: the characteristics of the APs.

4.4 Impact of AP Wealth on Prices

Let us explore how the profitability distribution F_v may affect the equilibrium prices. We still keep $\sigma_A : \sigma_B : \sigma_C = 1 : 3 : 5$ and

$(q_A, q_B, q_C) = (0.2, 1, 5)$. α_i follows the joint distribution of F_β and F_v , where F_β is distributed as *Geo*.




Figure 7: Shift in market prices as ρ varies: $(q_A, q_B, q_C) = (0.2, 1, 5)$, $\sigma_A : \sigma_B : \sigma_C = 1 : 3 : 5$, $F_\beta = \text{Geo}$.




Figure 8: Shift in market prices as ρ varies: $(q_A, q_B, q_C) = (0.2, 1, 5)$, $\sigma_A : \sigma_B : \sigma_C = 1 : 3 : 5$, $F_\beta = \text{Geo}$.

In Figure 7, we vary ρ from 0 to 1 along the x-axis and plot the equilibrium prices where the profitability distribution F_v follows a binomial distribution $BN(p)$ parameterized by $p = 0.2, 0.5$ and 0.8 respectively. By doing this, we simulate the normal distributions of the APs' wealth varying the mean value from small to large. We observe that despite the difference in mean profit of the APs, p_C drops to the minimum price at the same time. The price curves in all cases keep the same shape; however, they scale differently on the vertical axis. This indicates that the market prices depend on how much the APs are able to pay for the services, and how they demand for the TPs based on their values of (β_i, v_i) .

In Figure 8, we vary F_v to be *Geo*, *Uni* and *ReGeo*. We observe that the shapes of the price curves are very different: prices decrease convexly, linearly and concavely in the three subfigures. In general, how fast the prices drop depends on the density of the APs whose profitability are around that price range, and the shape of the curves look like the complimentary cumulative distribution function (CCDF) of F_v .

Lessons (the AP wealth effects on prices) leaned:

- The market prices of the TPs are positively correlated with the mean profitability of the APs.
- At a certain price range where the density of the APs is high (low), more (less) competition among the APs drives the prices close to (far below) their profitability.

4.5 Impact of AP Quality-Sensitivity on Prices

In this subsection, we study how the quality sensitivity distribution F_β affects the equilibrium prices. We set $\sigma_A : \sigma_B : \sigma_C = 1 : 3 : 5$ and $\rho = 0.5$. In the following cases, F_β follows a binomial distribution $BN(p)$, where we vary the parameter p along

the x-axis. By doing this, we simulate the cases where the APs become more and more sensitive to quality when the mean sensitivity increases with p .




Figure 9: Shift of market prices when we vary AP's sensitivity to quality: with $\sigma_A : \sigma_B : \sigma_C = 1 : 3 : 5$, $\rho = 0.5$, $(q_A, q_B, q_C) = (0.2, 1, 5)$ and $F_\beta = BN(p)$.




Figure 10: Shift of market prices when we vary AP's sensitivity to quality: with $\sigma_A : \sigma_B : \sigma_C = 1 : 3 : 5$, $\rho = 0.5$, $F_v = BN(0.5)$ and $F_\beta = BN(p)$.

In Figure 9, we fix $(q_A, q_B, q_C) = (0.2, 1, 5)$ and vary F_v to be *Uni*, *BN(0.5)* and *Geo* in the three sub-figures. We observe that although the profitability distribution affect the absolute price values, the shape of the price curves look similar. When the quality sensitivity of the APs increases, the lowest quality service price, i.e., p_C , drops sharply and quickly. Although p_A and p_B drops accordingly with p_C , after p_C reaches the minimum price, both p_A and p_B rebound. With further increase in quality sensitivity, p_B shows a trend to decrease slowly; however, p_A always stays at a high level. When the APs become more sensitive to quality, more and more APs start to move to higher class TPs. As a result, the capacity μ_C becomes under-utilized, which also drives p_C down very quickly. Although p_C 's drop pulls down the overall market prices, more APs move to higher class TPs, which make TPs *A* and *B* in demand, and therefore, keep p_A and p_B steadier. After p_C reaches the minimum price, p_C stops decreasing. As the APs' quality sensitivity keeps increasing, even the minimum market price of p_C becomes relatively expensive to the APs. This makes even more APs move to TP *A* and *B* and drives p_A and p_B upward.

In Figure 10, we fix $F_v = BN(0.5)$ and vary the qualities to be $(q_A, q_B, q_C) = \kappa(0.2, 1, 5)$, where $\kappa = 0.5, 1$ and 2 . We observe the same trends as in Figure 9 that p_C drops quickly and sharply to the minimum price as the APs' quality sensitivity increases. As κ increases, all the price curves move to the left and the price drop of lower class TPs becomes quicker and sharper. This also coincides with the observations made in Figure 5 that when the qualities degrade, the price of the lower class TP drops much quicker.

In the above illustrations, we vary the distribution F_β . It is also possible that all the APs' sensitivity increase by $\beta'_i = \xi \beta_i$ for some

$\xi > 1$. By Theorem 2, we can rescale the system by $\kappa_3 = 1/\xi$, as if the APs keep their quality sensitivity constant and all the qualities become poorer. By Theorem 3 and the TP quality effect result in Figure 5, we also conclude that more APs will prefer higher quality TPs and the price of the lower quality TPs will drop sharply.

Lessons (the AP quality-sensitivity effects on prices) leaned:

- When the APs become more sensitive to the service quality, the price of lower class TPs will drop quickly.
- When the price of the lowest quality TP goes down to its cost, the prices of higher quality TPs might increase due to their relatively cheap prices and high demand.

4.6 Internet Evolution: Some Explanations

By understanding the factors that drive the market equilibrium, we reason about the evolution of the Internet ecosystem and reach plausible answers to the questions raised in Section 1. We do not claim that our answers below are exhaustive and the limitations of our model will be discussed in Section 5.4.

1) Why have the IP transit prices been dropping? The capacity effect tells that the price drop can be a consequence of the capacity expansion of the transit providers. Compared to the capacities at the last-miles, the capacity in the backbone grows faster than demand and is abundant [14]. Also, the price drop in better quality services, i.e., CDN prices, will drive the transit prices further down. The quality effect tells that when the transit quality differs a lot from the CDN services, the prices will diverge greatly. The wealth effect tells that since the majority of the elastic APs might not be very profitable, transit providers cannot fully utilize its capacity and charge a high price at the same time. This is also why they are looking for providing value-added and differentiated services. Last, the AP quality-sensitivity effect tells that when AP traffic becomes more and more inelastic, e.g., the surge of Netflix traffic, lower quality service will become less valuable and therefore its price will drop quickly.

2) Why have the CDNs emerged in the ecosystem? The capacity effect tells that when the capacity of higher quality service is small, it can maintain a price difference with the lower quality services. The quality effect tells that if a CDN service's quality differs a lot from the transit services, it can be priced much higher. When the capacity of the transit market was limited and priced high, the demand for even higher quality service drove the price for potential CDN services even higher. This explains why CDNs emerged in the first place. The wealth effect tells that when the APs' profitability is not high, the market prices cannot be high. However, due to the low cost structure of the CDNs, they can still help small APs who could not afford the infrastructure to support large demand. The AP quality-sensitivity effect further tells that with the traffic being more and more sensitive to quality, the price of high quality CDN can sustain at a high level.

3) Why has the pricing power shifted to the access ISPs? This can be partially explained by the AP quality-sensitivity effect and the TP quality effect. When the AP traffic becomes more and more sensitive to service quality, they are more willing to pay for the higher quality services. Because access ISPs are physically closer to the users, their service quality is naturally much better than other providers who have to go through the access ISPs to reach the end-users anyways. Consequently, the difference in service quality makes it possible for the access ISPs to charge services at higher prices. Furthermore, Comcast's monopolistic position in the U.S. market could be another reason, under which its price will be set higher than the competitive market price under Definition 3.

4) Why are the large content providers building their own wide-area networks toward users? Mostly because the APs become more sensitive to service quality, they cannot rely on the transit providers to deliver content. As high quality services are limited and access ISPs would obtain more pricing power, large APs might consider establishing their own networks toward users as a cheaper alternative than paying access ISPs for better services in the future.

5. INTERNET'S ECONOMIC EVOLUTION

Besides understanding how each isolated factor might affect the market prices, we incorporate ground truth data [10, 19, 4, 11], e.g., the historical trends of the TPs' capacity expansion and the APs' characteristics, and project *possible future price dynamics* of the Internet ecosystem. Through this, our model can help the TPs make various long-term business decisions. Let us demonstrate this.

We take a macroscopic view and categorize network services as two types: $\mathcal{M} = \{A, B\}$. B models the IP transit service that provides interconnection based on "best-effort"; A models the CDN or private peering type of service that provides better service quality than B . We categorize the APs as three types: $\mathcal{N} = \{a, b, c\}$. a models the video or realtime interactive applications that are very sensitive to quality. b models the web applications that are elastic but more tolerate to quality than type a applications. c models the inelastic applications, e.g., email and P2P file download.

By Corollary 1, we know that when quality and the sensitivity parameters scale inversely, the equilibrium remains the same; therefore, without loss of generality, we set $q_B = 1$ as the baseline best-effort quality level. We set the quality sensitivity parameters to be $(\beta_a, \beta_b, \beta_c) = (10, 1, 0.1)$. Under this setting, type a APs would only obtain $e^{-10} \approx 4.5^{-5}$ of their maximum throughput under q_B , which implies that the best-effort service cannot support quality sensitive applications. Also, under q_B , a type b AP could get $e^{-1} \approx 37\%$ of its maximum throughput; however, a type c AP could get $e^{-0.1} \approx 90\%$ of its maximum throughput. When measured by delay, the quality of service for realtime applications often require the delay to be at the order of milliseconds [27], compared to the best-effort service delays at the order of seconds. Thus, we choose $q_A = 0.01$ to reflect the same order of magnitude of service difference. As a result, even type a APs would obtain $e^{-0.1} \approx 90\%$ of their maximum throughput under the better quality level q_A .

Next, we try to estimate the capacity of the TPs on the Internet. We take the Equinix Internet Exchange at New York (Equinix-NY) as a reference market and estimate the capacities based on the data provided by PeeringDB [10]. At the end of year 2011, there were 102 ISPs listed on at Equinix New York Exchange in PeeringDB, among which 44 use *Open* peering policy and the remaining 58 use either *Selective* or *Restricted* peering policy. The total capacity was around 21 Tbps, among which the ISPs using *Open* peering policy contributed 7 Tbps and the remaining ISPs contributed 14 Tbps. Since *Selective* and *Restricted* policies are used for private and often paid-peering agreements, we set ν_A and ν_B to be 14 and 7 Tbps, for the reference time of the year 2011.

From the Global Internet Geography [4] report, between 2007 to 2011, the international Internet capacity increased six-fold and the bandwidth to the U.S. had increased nearly 50 percent per year. To a first approximation, we assume that the capacity ν of the TPs increases 50% per year. We define $\alpha = \alpha_a + \alpha_b + \alpha_c$ and ω_a, ω_b and ω_c as the weight of the throughput upper bound of each application type. Given α and the weight of AP i , we obtain α_i as

$$\alpha_i = \frac{\omega_i}{\omega_a + \omega_b + \omega_c} \alpha, \quad \forall i = a, b, c.$$

Based on the observed traffic distribution of various applications in [19], we set $(\omega_a, \omega_b, \omega_c) = (2\%, 75\%, 23\%)$ for the year 2007, and assume that the weight for video (ω_a), web (ω_b) and inelastic applications (ω_c) increase at an annual growth rate of 150%, 50% and 20% respectively. Notice that IP transit prices are often quoted for per Mbps-month, while CDN prices are often quoted for per terabit. If capacity is fully utilized 24/7, \$1 per Mbps-month can be translated into \$0.386 per terabit. We assume that the maximum per unit traffic revenue for the APs is \$10 Mbps-month and the APs' revenue are uniformly distributed.

5.1 A First Approximation Benchmark

We use our macroscopic model to fit the historical prices starting from 2007 and project future Internet prices. In a first approximation, we choose the following parameters.

1. α at year 2007 (denoted as α_{07}) equals 10 Tbps.
2. α increases at an annual growth rate $r_\alpha = 22\%$.
3. $\eta_A = \mu_A/\nu_A = 0.3$ and $\eta_B = \mu_B/\nu_B = 0.9$.




Figure 11: Historical price and future price projection.

In Figure 11, the upper left subfigure plots the achievable throughput for the CDN (μ_A) and IP transit (μ_B) services from 2007 to 2014 and the lower left subfigure plots the maximum demand α_a , α_b and α_c for the same time period. The upper right subfigure plots the price dynamics of both IP transit and CDN services and the lower right subfigure plots the percentage of price change for both services. We observe that average price drop from 2007 to 2011 is approximately 20%, which coincides with the price drop surveyed in the Global Internet Geography [4] report. Also, the price of IP transit is below \$2 per Mbps-month, very close to the mean of IP transit prices, where the lowest price fell to \$1 per Mbps-month.

Compared to the video delivery pricing [11], our price projection shows that the CDN price drops around 8% annually from 2007 to 2011, and reaches \$5.67 per Mbps-month, or \$2.18 per terabit. This price is lower than the \$7.5 per terabit price for APs with volume of 5PB data and the price drop is slower than the observed 20% price drop in the CDN industry [11]. The difference could come for two reasons: 1) since CDN service charges based on traffic volume, we cannot assume that the APs would always use the capacity 24/7, and therefore, the CDN providers should charge some premium on top of the basic per Mbps per month charge, 2) in contrast to our competitive model for CDN service, the industry

might be less competitive and could charge a much higher price; therefore, when the industry becomes more competitive, we expect to see much sharper price drops.

Based on the trend from 2007 to 2011, our model projects that both the IP transit and CDN prices will further drop, at an even faster rate, and IP transit price will drop to its minimum price. Of course, this projection is based on the assumption that the capacity of the TPs will keep expanding at the 50% annual rate. We will further discuss potential trends of future prices in a later subsection.

5.2 Sensitivity of the Benchmark

In this subsection, we show the sensitivity of our price projection with respect to the chosen parameters.




Figure 12: Sensitivity to initial demand α_{07} and rate r_α .

First, we want to see how the demand parameter α affects the price dynamics. Figure 12 shows a projection of service prices when the initial value α_{07} and the growth rate r_α change. In the left subfigure, we vary α_{07} to be 9 and 11 Tbps compared to the benchmark value of 10 Tbps. We observe that the prices are positively correlated with α_{07} . In the right subfigure, we vary the growth rate r_α to be 20% and 24% compared to the benchmark value of 22%. We observe that the prices again are positively correlated with r_α . Both tells that when the demand increases, so do the prices.




Figure 13: Sensitivity to capacity utilization η_A and η_B .

Second, we want to see how the capacity utilization factor $\eta_I = \mu_I/\nu_I$ affects the price dynamics. Figure 13 shows a projection of Internet service prices when η_A and η_B vary from the benchmark. In the left subfigure, we vary η_A to be 0.2 and 0.4 compared to the benchmark value 0.3. In the right subfigure, we vary η_B to be 0.8 and 1.0 compared to the benchmark value 0.9. We observe that the all prices are negatively correlated with the capacity utilization factors. Also, the IP transit prices are sensitive to both η_A and η_B ; while the CDN prices are only sensitive to its utilization factor η_A .

5.3 Price Projection and TP Business Decisions

Now, we demonstrate that by using the price projection from our model, we can help the TPs to make business decisions on 1) how aggressive the TPs should expand their capacity, and 2) whether the TPs should/would tend towards *Open* or *Selective* peering policies.




Figure 14: Price projection under various capacity ratios $\nu_A : \nu_B$ and capacity expansion rate r_ν .

To answer the first question, we vary the capacity growth rate of the TPs r_ν to be 40% and 60%, compared to the historical benchmark rate 50% and plot the price projections in the left subfigure of Figure 14. We observe that when the capacity grows at 60% per year, both the CDN and IP transit price drop fast and the IP transit price will down to its cost next year; however, when the growth rate is 40%, the IP transit price will be decreasing at a very slow rate. These observations tell us that the ISPs providing IP transit services might want to slow down their investment in capacity expansion; however, CDN providers and ISPs that sell private-peering and QoS might want to continue to expand their capacity when their profit margins are still above zero. As the price of IP transit drops, we believe that the investment in the transit capacity will slow down, which will also stabilize the price of the IP transit services.

To answer the second question, we vary the capacity ratio $\nu_A : \nu_B$ from the benchmark ratio 2 : 1 (14 Tpbs : 7 Tpbs) to 3 : 1 and 3 : 2 for the year 2014 in the right subfigure. These two projections model the scenarios where ISPs will tend to be more *Selective* and more *Open* in their peering policies respectively. We observe that if more ISPs are going to use an *Open* peering policy, the IP transit price will drop to its cost quickly; otherwise, the IP transit price will get closer to the CDN price and be stable. This observation implies that ISPs would have strong incentives to move towards *Selective* peering policies if possible, which coincides with the reality that the access ISP, Comcast, started to use private peering exclusively.

In summary, we predict that although the CDN price will still be dropping, the price of IP transit will be more stable. Furthermore, the capacity expansion will slow down and more ISPs will tend to use *Selective* rather than *Open* peering policies in the near future.

5.4 Limitations of the Model

First of all, our general equilibrium model implicitly assumes that each market segment is competitive. In practice, some market segment could be lack of competition and form a monopolistic or oligopolistic market structure. Thus, the real market prices will be higher than what our model predicts. Second, our equilibrium model does not capture the off-equilibrium and transit dynamics that could happen in practice. Third, our model is in nature macroscopic, and it does not capture detailed information like peering agreement, topology, traffic patterns and etc. Nevertheless, our model does capture the type of different services the TPs provide

via implicitly encoding all the relevant information into the quality level q_I . From the APs' point of view, they do care about *quality* rather than other details of the TPs. Fourth, since our focus is on the transit/CDN market, our model does not intend to capture the end-user market aspects. For example, modeling the bundle of access services and other service differentiations are out of scope. Last but not the least, our macroscopic model provides some qualitative reasons for the Internet evolution, which we do not claim to be exhaustive. There might be additional factors/reasons that are not captured by our model, e.g., the lack of competition in the market.

6. RELATED WORK

Many empirical studies have been tracking the evolution of the Internet using measurements and public data sets [19, 15, 17, 24, 12]. Labovitz et al. [19] 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. [17] collected and analyzed traceroute measurements and showed that large content providers are deploying their own wide-area networks. Dhamdhere et al. [15] confirmed the consolidation of the core of the Internet, that brings the content closer to users. Akella et al. [12] 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. [24] studied the network access link performance measured directly from home gateway devices. We focus on a macroscopic model of the Internet ecosystem that captures the application traffic going through the network transport service providers.

Many works [13, 16, 21, 22, 25, 20, 18] focused on the modeling perspective of the Internet evolution. Chang et al. [13] presents an evolutionary model for the AS topologies. Lodhi et al. [20] used an agent-based model to study the network formation of the Internet. Motiwala et al. [22] used a cost model to study the Internet traffic. Valancius et al. combined models and data to study the pricing [25] structure of the IP transit market. Faratin et al. [16] and Ma et al. [21] studied the evolution of the ISP settlements. In this work, we take a holistic view and analyze the business decisions and evolutions of the APs and TPs altogether.

7. CONCLUSIONS

We proposed a network aware, macroscopic model to explain the evolution of the Internet. Our model captures 1) the business decisions of the APs, 2) the pricing and competition of the TPs, and 3) the resulting market equilibrium of the ecosystem. By analyzing how the AP characteristics (i.e., traffic intensity, profitability and sensitivity to quality), and the TP characteristics (i.e., quality, price and capacity, affect the market equilibrium), we obtain fundamental understanding of why historical and recent evolutions of the Internet have happened. With further estimations of the trends in traffic demand, capacity growth and quality improvements, our model can also project the future evolution of the Internet ecosystem. This model provides a tool for the Internet players to better understand their business and risks, and help them to deal with their business decisions in the complicated and evolving ecosystem.

Acknowledgement

This study is supported by the Human Sixth Sense Programme at the Advanced Digital Sciences Center from Singapore's Agency for Science, Technology and Research (A*STAR), Ministry of Education of Singapore AcRF grant R-252-000-448-133, and the National Science Foundation grants CNS-1017934 and CCF-1139915.

8. REFERENCES

- [1] Akamai. <http://www.akamai.com/>.
- [2] Amazon Elastic Compute Cloud (EC2). <http://www.amazon.com/ec2>.
- [3] Cogent Communications, Inc. <http://www.cogentco.com>.
- [4] "Global Internet Geography." Telegeography Research. <http://www.telegeography.com/>.
- [5] Level 3 Communications, Inc. <http://www.level3.com>.
- [6] Limelight Networks. <http://www.limelight.com/>.
- [7] Netflix, Inc. <http://www.netflix.com>.
- [8] Netflix takes up 32.7% of Internet bandwidth, CNN News. <http://edition.cnn.com/2011/10/27/tech/web/netflix-internet-bandwidth-mashable>.
- [9] Netflix Traffic Now Bigger Than BitTorrent. Has Hollywood Won? GIGAOM News. <http://gigaom.com/broadband/netflix-p2p-traffic/>.
- [10] PeeringDB. <http://www.peeringdb.com/>.
- [11] Video Delivery Pricing for Q4 2011. <http://www.cdnpricing.com/>.
- [12] 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.
- [13] 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.
- [14] C. Courcoubetis and R. Weber. *Pricing Communication Networks: Economics, Technology and Modelling*. John Wiley & Sons Ltd., 2003.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] H. Haddadi, S. Uhlig, A. Moore, R. Mortier, and M. Rio. Modeling Internet topology dynamics. *ACM SIGCOMM Computer Communication Review, Volume 38 Issue 2, April 2008*.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] W. Norton. *The Internet Peering Playbook: Connecting to the Core of the Internet*. DrPeering Press, 2011.
- [24] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapill. Broadband Internet performance: A view from the gateway. In *Proceedings of ACM SIGCOMM*, Toronto, Ontario, Canada, August 2011.
- [25] 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.
- [26] T. Wu. Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law*, 141, 2005.
- [27] X. Xiao. *Technical, Commercial and Regulatory Challenges of QoS: An Internet Service Model Perspective*. The Morgan Kaufmann Series in Networking, 2008.