

Congestion Equilibrium for Differentiated Service Classes

Richard T. B. Ma

Advanced Digital Sciences Center, Illinois at
Singapore and National University of Singapore

Vishal Misra

Columbia University

ABSTRACT

In this paper, we introduce the notation of a *congestion equilibrium* where competing players share a common level of congestion in the system. We illustrate the concept by developing a framework under which content providers (CPs) compete for a last-mile bottleneck capacity of an ISP so as to service their users. We show that under minor monotonicity assumptions on the system congestion and traffic rates, a unique congestion equilibrium exists. Under an independence of scale assumption, we characterize the congestion and rates in equilibrium as continuous homogeneous functions. In particular, the congestion in equilibrium can be characterized by a continuous non-increasing function of the per capita capacity of the system. Through examples, we show that various congestion metrics apply to our model, e.g. $M/M/1$ delay metric and throughput metrics that result from work-conserving congestion control mechanism.

By using the congestion equilibrium as a building block, we further model class-based service differentiations of an ISP under which CPs choose service class and compete resource with other CPs in the same service class. We formulate a two-stage strategic game between an ISP and the CPs. We characterize the Nash equilibrium and competitive equilibrium, which is based on a corresponding *congestion-taking assumption*, and propose an algorithm to practical solve a competitive equilibrium.

1. CONGESTION EQUILIBRIUM MODEL

We consider a model of the Internet with three parties: 1) CPs, 2) ISPs and 3) consumers. We focus on a fixed consumer group in a targeted geographic region. We denote M as the number of consumers in the region¹. Each consumer subscribes to an Internet access service via an ISP. We consider the scenarios where one monopolistic ISP I or a set \mathcal{I} of competing oligopolistic ISPs provide the Internet access for the consumers. We denote \mathcal{N} as the set

¹Note that M can also be interpreted as the *average* or *peak* number of consumers accessing the Internet simultaneously in the region, which will scale with the total number of actual consumers in a region. This does not change the nature of any of the results we describe subsequently, but gives a more realistic interpretation of the rate equilibrium.

of CPs from which the consumers request content. We define $N = |\mathcal{N}|$ as the number of CPs. Our model does not include the backbone ISPs for two reasons. First, the bottleneck of the Internet is often at the last-mile connection towards the consumers [6], both wired and wireless. We focus on the regional or so-called eyeball ISPs [7] that provide the bottleneck last-mile towards the consumers. Second, the recent concern on network neutrality manifests itself in the cases where the last-mile ISPs, e.g. France Telecom, Telecom Italia and Vodafone, intended to differentiate services and charge CPs, e.g. Apple and Google, for service fees [3].

We denote μ as the last-mile bottleneck capacity towards the consumers in the region. Figure 1 depicts the contention




Figure 1: Contention at the last-mile bottleneck link.

at the bottleneck among different flows from the CPs. We denote λ_i as the aggregate throughput rate from CP i to the consumers. Because consumers initiate downloads and retrieve content from the CPs, we first model the consumer demands so as to characterize the CPs' throughput rates λ_i s. Given a set \mathcal{N} of CPs, a group of M consumers and a link with capacity μ , we denote the system as a triple (M, μ, \mathcal{N}) .

1.1 Rate as a Function of Congestion

We denote $\hat{\theta}_i$ as the *unconstrained throughput* for a typical user of CP i . For instance, the unconstrained throughput for the highest quality Netflix streaming movie is about 5 Mbps [2], and given an average query page of 20 KB and an average query response time of .25 seconds [1], the unconstrained throughput for a Google search is about 600 Kbps, or just over $1/10^{th}$ of Netflix. We denote $\alpha_i \in (0, 1]$ as the percentage of consumers that ever access CP i 's content, which models the popularity of the content of CP i . We define $\hat{\lambda}_i = \alpha_i M \hat{\theta}_i$ as the unconstrained throughput of CP i . We denote λ_i as the rate of CP i and define the aggregate rate as $\lambda_{\mathcal{N}} = \sum_{i \in \mathcal{N}} \lambda_i$. Without contention, CP i 's

throughput λ_i equals $\hat{\lambda}_i$. However, when the capacity μ cannot support the unconstrained throughput from all CPs, e.g. $\mu < \sum_{i \in \mathcal{N}} \hat{\lambda}_i$, each CP's rate $\lambda_i \leq \hat{\lambda}_i$. In general, CP i 's rate λ_i depends on the level of congestion in the network. We denote a non-negative real number Γ as the level of congestion caused by the set of competing CPs \mathcal{N} . We will see that Γ can represent various congestion metrics in general. Each CP's rate can be expressed as a function² of Γ as

$$\lambda_i(M, \mu, \mathcal{N}) = \lambda_i(\Gamma) = \alpha_i M \rho_i(\Gamma),$$

where $\rho_i(\Gamma)$ can be interpreted as the per-user achievable rate under congestion Γ .

Assumption 1. For any CP i , $\rho_i(\cdot)$ is a non-negative, continuous and non-increasing function defined on the domain of $[0, \hat{\theta}_i]$, and satisfies $\rho_i(0) = \hat{\theta}_i$ and $\lim_{\Gamma \rightarrow \infty} \rho_i(\Gamma) = 0$.

1.2 Congestion as a Function of Throughput Rates and Capacity

When multiple flows share the same bottleneck link, they compete for capacity and therefore, induce a level of congestion in the system. We can also view the congestion Γ as a function of rates $\Lambda = (\lambda_1, \dots, \lambda_N)$ of the CPs and the system capacity μ .

Assumption 2 (MONOTONICITY). $\Gamma : \mathbb{R}^{|\mathcal{N}|} \times \mathbb{R} \rightarrow \mathbb{R}$ is a continuous function. For any capacity $\mu_1 \leq \mu_2$ and rates $\Lambda' \leq \Lambda''$, the system congestion $\Gamma(\cdot)$ satisfies

$$\Gamma(\Lambda, \mu_1) \geq \Gamma(\Lambda, \mu_2) \quad \forall \Lambda \geq 0, \quad \text{and}$$

$$\Gamma(\Lambda', \mu) \leq \Gamma(\Lambda'', \mu) \quad \forall \mu \geq 0.$$

Moreover, given a fixed capacity μ , the minimum congestion $\Gamma_{\mathcal{N}}(\mathbf{0}, \mu)$ is upper-bounded.

Assumption 3 (INDEPENDENCE OF SCALE). For any $\xi > 0$, the system congestion Γ satisfies

$$\Gamma(\Lambda, \mu) = \Gamma(\xi\Lambda, \xi\mu).$$

The Independence of Scale assumption states that if the capacity scales at the same rate as the consumer size, the level of system congestion Γ remains the same.

1.3 Rates and Congestion in Equilibrium

The congestion functions map the throughput rates to a level of congestion; the rate functions specify rate based on a given level of congestion. The interplay between a congestion function and the rate functions determines the system congestion and rates in equilibrium.

Definition 1. A pair (Λ, Γ) is a congestion equilibrium of the system (M, μ, \mathcal{N}) if $\lambda_i = \lambda_i(M, \mu, \mathcal{N}) = \alpha_i M \rho_i(\Gamma)$ for all $i \in \mathcal{N}$ and $\Gamma = \Gamma(\Lambda, \mu)$.

²We use λ_i as a fixed rate and $\lambda_i(\cdot)$ as a function. Similarly, we use Γ and Λ as fixed level of congestion and rates, and $\Gamma(\cdot)$ and $\Lambda(\cdot)$ as functions.

Theorem 1. Under Assumption 1 and 2, a system (M, μ, \mathcal{N}) has a unique congestion equilibrium (Λ, Γ) .

Proof of Theorem 1: We first show the existence of an equilibrium. Let $\hat{\Lambda} = (\hat{\lambda}_1, \dots, \hat{\lambda}_N)$. If $\Gamma(\hat{\Lambda}, \nu) = 0$, then $(\Lambda, \Gamma) = (\hat{\Lambda}, 0)$ is an equilibrium of the system. We consider the cases where $\Gamma(\hat{\Lambda}, \nu) > 0$. Because $\rho_i(\cdot)$'s are non-increasing functions, $\lambda_i(\gamma) = \alpha_i M \rho_i(\gamma)$ are non-increasing functions of γ . We construct a vector parameterized by γ as $\Lambda(\gamma) = (\lambda_1(\gamma), \dots, \lambda_N(\gamma))$. Therefore, $\Lambda(\gamma)$ are non-increasing in γ the \mathbb{R}^N space. By the monotonicity property of Assumption 2, $\Gamma(\Lambda(\gamma), \mu)$ is a non-increasing function of γ . Since $\lim_{\gamma \rightarrow \infty} \rho_i(\gamma) = 0$, $\lim_{\gamma \rightarrow \infty} \Gamma(\Lambda(\gamma), \mu) = \Gamma(\mathbf{0}, \mu)$, which is upper-bounded. Since $\Lambda(0) = \hat{\Lambda}$ and $\Gamma(\Lambda(0), \mu) > 0$, there exists a solution γ^* that satisfies

$$\Gamma(\Lambda(\gamma^*), \mu) = \gamma^*.$$

$(\Lambda, \Gamma) = (\Lambda(\gamma^*), \gamma^*)$ is an equilibrium of the system.

Next, we show the uniqueness of the equilibrium by contradiction. Suppose we have two different equilibria (Λ', Γ_1) and (Λ'', Γ_2) . If $\Gamma_1 = \Gamma_2$, $\Lambda' = \Lambda(\Gamma_1) = \Lambda(\Gamma_2) = \Lambda''$, which implies two equal equilibria. Without loss of generality, we assume $\Gamma_1 < \Gamma_2$. By Assumption 1, we have $\Lambda(\Gamma_1) \geq \Lambda(\Gamma_2)$. By Assumption 2, we have $\Gamma(\Lambda(\Gamma_1), \mu) \geq \Gamma(\Lambda(\Gamma_2), \mu)$. However, this implies $\Gamma_1 \geq \Gamma_2$, which contradicts our assumption. ■

Without confusion, we specify an equilibrium of a system as its unique congestion equilibrium. By Assumption 3, we further characterize the equilibrium (Λ, Γ) as follows.

Theorem 2. Under Assumption 1 to 3, if (Λ, Γ) is an equilibrium of the system (M, μ, \mathcal{N}) , then for any $\xi > 0$, $(\xi\Lambda, \Gamma)$ is the unique equilibrium of the system $(\xi M, \xi\mu, \mathcal{N})$.

Proof of Theorem 2: Let (Λ', Γ') be an equilibrium of the system $(\xi M, \xi\mu, \mathcal{N})$. In the new system, with rates $\Lambda' = \xi\Lambda$, the induced level of congestion is $\Gamma(\xi\Lambda, \xi\mu)$. By Assumption 3, $\Gamma(\xi\Lambda, \xi\mu) = \Gamma(\Lambda, \mu) = \Gamma$. This implies that in the new system $\lambda'_i = \alpha_i \xi M \rho_i(\Gamma) = \xi \lambda_i$ and $\Lambda' = \xi\Lambda$. Therefore, $(\Lambda', \Gamma') = (\xi\Lambda, \Gamma)$ is an equilibrium of the system $(\xi M, \xi\mu, \mathcal{N})$. By Theorem 1, it is also the unique equilibrium of the system. ■

We define $\nu = \mu/M$ as the per capita capacity of the system. Given a fixed set of CPs \mathcal{N} , we denote $\Gamma_{\mathcal{N}}(M, \mu) = \Gamma(M, \mu, \mathcal{N})$ and $\Lambda_{\mathcal{N}}(M, \mu) = \Lambda(M, \mu, \mathcal{N})$ as the unique level of congestion and rates in equilibrium.

Corollary 1. The system congestion $\Gamma_{\mathcal{N}}(M, \mu)$ is a homogeneous function of degree 0, i.e.

$$\Gamma_{\mathcal{N}}(\nu) = \Gamma_{\mathcal{N}}(\xi M, \xi\mu) \quad \forall \xi > 0.$$

Moreover, $\Gamma_{\mathcal{N}}(\nu)$ is a continuous non-increasing function of the per capita capacity ν and for all $\nu > 0$, it satisfies

$$\Gamma_{\mathcal{N}_1}(\nu) \leq \Gamma_{\mathcal{N}_2}(\nu) \quad \forall \mathcal{N}_1 \subseteq \mathcal{N}_2.$$

Proof of Corollary 1: As a result of Theorem 2, linearly scaled systems will induce the same level of congestion in

equilibrium; therefore, $\Gamma_{\mathcal{N}}$ can be expressed as a function of $\nu = \mu/M$. To show that $\Gamma_{\mathcal{N}}(\nu)$ is non-increasing in ν , we want to show $\Gamma_{\mathcal{N}}(\nu_1) \geq \Gamma_{\mathcal{N}}(\nu_2)$ for all $\nu_1 \leq \nu_2$. Because the system congestion does not change in a linearly scaled system, we only need to show the case where $M_1 = M_2 = M$ and $\mu_1 \leq \mu_2$. We prove by contradiction and suppose $\Gamma_{\mathcal{N}}(\nu_1) < \Gamma_{\mathcal{N}}(\nu_2)$ which implies $\Lambda(\Gamma_{\mathcal{N}}(\nu_1)) \geq \Lambda(\Gamma_{\mathcal{N}}(\nu_2))$ by Assumption 1. By Assumption 2, we further have $\Gamma(\Lambda(\Gamma_{\mathcal{N}}(\nu_1)), \mu_1) \geq \Gamma(\Lambda(\Gamma_{\mathcal{N}}(\nu_2)), \mu_2)$ which is equivalent to $\Gamma_{\mathcal{N}}(\nu_1) \geq \Gamma_{\mathcal{N}}(\nu_2)$. Therefore, we have a contradiction.

Similarly, we show for any $\mathcal{N}_1 \subseteq \mathcal{N}_2$, $\Gamma_{\mathcal{N}_1}(\nu) \leq \Gamma_{\mathcal{N}_2}(\nu)$ by contradiction. Suppose $\Gamma_{\mathcal{N}_1}(\nu) > \Gamma_{\mathcal{N}_2}(\nu)$, by Assumption 1, we have $\lambda_i(\Gamma_{\mathcal{N}_1}(\nu)) \leq \lambda_i(\Gamma_{\mathcal{N}_2}(\nu))$ for all $i \in \mathcal{N}_1$. We create an augmented rate vector Λ' , where $\lambda'_i = \lambda_i(\Gamma_{\mathcal{N}_1}(\nu))$ for all $i \in \mathcal{N}_1$ and virtually assign $\lambda'_i = 0$ for all $i \in \mathcal{N}_2 \setminus \mathcal{N}_1$. Thus, we have $\Lambda' \leq \Lambda(\Gamma_{\mathcal{N}_2}(\nu))$. By Assumption 2, we further have $\Gamma(\Lambda(\Gamma_{\mathcal{N}_1}(\nu)), \mu) = \Gamma(\Lambda', \mu) \leq \Gamma(\Lambda(\Gamma_{\mathcal{N}_2}(\nu)), \mu)$. This is equivalent to $\Gamma_{\mathcal{N}_1}(\nu) \leq \Gamma_{\mathcal{N}_2}(\nu)$, which shows a contradiction. ■

Corollary 2. *The system rates $\Lambda_{\mathcal{N}}(M, \mu)$ is a homogenous function of degree -1 , i.e.*

$$\Lambda_{\mathcal{N}}(M, \mu) = \xi^{-1} \Lambda_{\mathcal{N}}(\xi M, \xi \mu) \quad \forall \xi > 0.$$

Proof of Corollary 2: By Theorem 2, any linearly scaled system induces the same level of congestion. We have

$$\lambda_i(\xi M, \xi \mu, \mathcal{N}) = \alpha_i \xi M \rho_i(\Gamma) = \xi \lambda_i(M, \mu, \mathcal{N}).$$

Therefore, $\Lambda_{\mathcal{N}}(\xi M, \xi \mu) = \xi \Lambda_{\mathcal{N}}(M, \mu)$. ■

1.4 Congestion Metrics and Systems

In this subsection, we illustrate various system mechanisms and their corresponding congestion metrics that can be modeled by our congestion equilibrium. Before presenting more detailed examples, we mention that there is a natural decomposition of the per-user achievable rate $\rho_i(\Gamma)$ as

$$\rho_i(\Gamma) = d_i(\Gamma) \theta_i(\Gamma), \quad (1)$$

where $\theta_i(\Gamma)$ is the achievable rate of an active user and $d_i(\Gamma)$ is percentage of users still being active under congestion Γ .

1.4.1 FIFO with delay as a congestion metric

In the simplistic example, the bottleneck ISP only perform a first-come-first-serve service discipline on the packet flows. Since all data packets encounter the same queueing delay under the FIFO scheduler, we can naturally measure the level of congestion by the delay of the system. Under a M/M/1 model, the congestion function becomes

$$\Gamma(\Lambda, \mu) = \Gamma_{\mathcal{N}} = \frac{1}{\mu - \lambda_{\mathcal{N}}}.$$

1.4.2 Rate control and proportional share mechanism

Besides delay-based congestion metrics, achieved throughput is another measure of congestion. Suppose the bottleneck ISP perform a rate control mechanism that assign throughput for flows proportional their maximum demand, i.e.

$$\theta_i : \theta_j = \hat{\theta}_i : \hat{\theta}_j, \quad \forall i \in \mathcal{N}.$$

Under this system mechanism, we can use the following system-wide congestion metric:

$$\Gamma_{\mathcal{N}} = \frac{\hat{\theta}_i}{\theta_i} - 1, \quad \forall i \in \mathcal{N}.$$

Moreover, the per-user throughput of CP i could be expressed as

$$\theta_i = \frac{\hat{\theta}_i}{\Gamma_{\mathcal{N}} + 1}, \quad \forall i \in \mathcal{N}.$$

1.4.3 End-to-end congestion control mechanism

Due to the end-to-end design principle of the Internet [5], congestion control has been implemented by window-based protocols, i.e. TCP and its variations. Mo and Walrand [10] showed that a class of α -proportional fair solutions can be implemented by window-based end-to-end protocols. Among the class of α -proportional fair solutions, the max-min fair allocation, a special case with $\alpha = \infty$, is the result of the AIMD mechanism of TCP [4]. Differing round trip times, receiver window sizes and loss rates can result in different bandwidths, but to a first approximation, TCP provides a max-min fair allocation of available bandwidth amongst flows.

Given an end-to-end congestion mechanism implemented, the system-wide congestion metric could be defined as

$$\Gamma_{\mathcal{N}} = \frac{1}{\max\{\theta_i : i \in \mathcal{N}\}}.$$

Similarly, the per-user throughput of CP i could be expressed as

$$\theta_i = \max\{\hat{\theta}_i, \frac{1}{\Gamma_{\mathcal{N}}}\}.$$

Notice that for all the three examples, the congestion function does not have to be unique; however, all of them satisfy the monotonicity and independent of scale properties.

2. DIFFERENTIATED SERVICE CLASSES

We assume that the monopolistic last-mile ISP I has a capacity of μ . This ISP can be a retail residential ISP, e.g. Comcast and Time Warner Cable, or a mobile operator, e.g. Verizon and AT&T. Regardless of whether it is a wired or wireless provider, it serves as the last-mile service provider for the consumers. We assume that the ISP is allowed to allocate a fraction $\kappa \in [0, 1]$ of its capacity to serve premium CPs and charge them at a rate $c \in [0, \infty)$ (dollar per unit traffic). For a wired ISP, κ can be interpreted as the percentage of capacity deployed for private peering points that charge a fee of c per unit incoming traffic and $1 - \kappa$ can be

interpreted as the percentage of capacity deployed for public peering points where incoming traffic is charge-free. For a wireless ISP, κ can be interpreted as the percentage of capacity devoted for the premium traffic that will be charged at a rate of c . The pair of parameters (κ, c) can also be thought of a type of Paris Metro Pricing (PMP) [11, 13], where one ordinary and another premium service class have capacities of $(1 - \kappa)\mu$ and $\kappa\mu$ and charge 0 and c respectively. In reality, content might be delegated via content distribution networks (CDNs), e.g. Akamai, or backbone ISPs, e.g. Level3 is a major tier-1 ISP that delivers Netflix traffic towards regional ISPs. Therefore, in practice, the charge c might be imposed on the delivering ISP, e.g. Level3, and then be recouped from the CP, e.g. Netflix, by its delivering ISP, e.g. Level3. Our model does not assume any form of the implementation.

We denote \mathcal{O} and \mathcal{P} as the two disjoint sets of CPs that join the ordinary and premium class respectively. We denote v_i as CP i 's per unit traffic revenue. This revenue can be generated by advertising for media clients, e.g Google, or by selling products to online consumers, e.g. Amazon, or by providing services to consumers, e.g. Netflix and e-banking. Our model does not assume how the revenue is generated either. Each CP i 's utility function u_i can be expressed as

$$u_i(\lambda_i) = \begin{cases} v_i \lambda_i & \text{if } i \in \mathcal{O}, \\ v_i \lambda_i - c \lambda_i & \text{if } i \in \mathcal{P}. \end{cases} \quad (2)$$

2.1 Content Provider's Best Response

Given the ISP's decision κ and c , each CP chooses whether to join the ordinary service class \mathcal{O} or the premium class \mathcal{P} .

Lemma 1. *Given a fixed set \mathcal{O} of CPs in the ordinary class and a fixed set \mathcal{P} of CPs in the premium class, a new CP i 's optimal strategy is to join the premium service class, if*

$$(v_i - c) \rho_i(\Gamma_{\mathcal{P} \cup \{i\}}(\kappa\nu)) \geq v_i \rho_i(\Gamma_{\mathcal{O} \cup \{i\}}((1 - \kappa)\nu)). \quad (3)$$

Moreover, when equality reaches, CP i obtains the same utility in both service classes.

Proof of Lemma 1: The utility of CP i joining the premium and ordinary service classes are $(v_i - c) \lambda_i(\Gamma_{\mathcal{P} \cup \{i\}}(\kappa\nu))$ and $v_i \lambda_i(\Gamma_{\mathcal{O} \cup \{i\}}((1 - \kappa)\nu))$ respectively. By dividing the constant $\alpha_i M$ on both, we obtain the above condition. ■

Lemma 1 states that a CP will join the premium service class if that results higher profit, which is per-unit flow profit $(v_i - c$ for the premium class) multiplied by the per capita throughput ρ_i . The above decision is clear for a CP only if all other CPs have already made their choices. To treat all CPs equally, we model the decisions of all CPs as a simultaneous-move game as part of a two-stage game.

2.2 Two-Stage Strategic Game

We model the strategic behavior of the ISP and the CPs as a two-stage game, denoted as a quadruple (M, μ, \mathcal{N}, I) .

1. *Players:* The ISP I and the set of CPs \mathcal{N} .

2. *Strategies:* ISP I chooses a strategy $s_I = (\kappa, c)$. Each CP i chooses a binary strategy of whether to join the premium class. The CPs' strategy profile can be written as $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$, where $\mathcal{O} \cup \mathcal{P} = \mathcal{N}$ and $\mathcal{O} \cap \mathcal{P} = \emptyset$.
3. *Rules:* In the first stage, ISP I decides $s_I = (\kappa, c)$ and announces it to all the CPs. In the second stage, all the CPs make their binary decisions simultaneously and reach a joint decision $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$.
4. *Outcome:* The set \mathcal{P} of the CPs shares a capacity of $\kappa\mu$ and the set \mathcal{O} of the CPs shares a capacity of $(1 - \kappa)\mu$. Each CP $i \in \mathcal{O}$ gets a rate $\lambda_i(M, (1 - \kappa)\mu, \mathcal{O})$ and each CP $j \in \mathcal{P}$ gets a rate $\lambda_j(M, \kappa\mu, \mathcal{P})$.
5. *Payoffs:* Each CP i 's payoff is defined by the utility $u_i(\lambda_i)$ in Equation (2). The ISP's payoff is the revenue $c\lambda_{\mathcal{P}}$ received from the premium class.

If we regard the set of CPs as a single player that chooses a strategy $s_{\mathcal{N}}$, our two-stage game can be thought of a Stackelberg game [12]. In this game, the first-mover ISP can take all the best-responses of the CPs into consideration and derive its optimal strategy s_I using *backward induction* [9]. Given any fixed strategy $s_I = (\kappa, c)$, the CPs derive their best strategies under a simultaneous-move game, denoted as $(M, \mu, \mathcal{N}, s_I)$. We denote $s_{\mathcal{N}}(M, \mu, \mathcal{N}, s_I) = (\mathcal{O}, \mathcal{P})$ as a strategy profile of the CPs under the game $(M, \mu, \mathcal{N}, s_I)$. Technically speaking, when $\kappa = 0$ or 1, there is only one service class. When $\kappa = 0$, we define the trivial strategy profile as $s_{\mathcal{N}} = (\mathcal{N}, \emptyset)$; when $\kappa = 1$, although there is not a physical ordinary class, we define the trivial strategy profile as $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{N} \setminus \mathcal{O})$, with $\mathcal{O} = \{i : v_i \leq c, i \in \mathcal{N}\}$ which defines the set of ISPs that cannot afford to join the premium class. Based on Lemma 1, we can define an equilibrium in the sense of a Nash or competitive equilibrium. To break a tie, we assume that a CP always chooses to join the ordinary service class when both classes provide the same utility.

Definition 2. *A strategy profile $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$ is a Nash equilibrium of a game $(M, \mu, \mathcal{N}, s_I)$, if*

$$\frac{v_i - c}{v_i} \begin{cases} \leq \frac{\rho_i(\Gamma_{\mathcal{O}}((1 - \kappa)\nu))}{\rho_i(\Gamma_{\mathcal{P} \cup \{i\}}(\kappa\nu))} & \text{if } i \in \mathcal{O}, \\ > \frac{\rho_i(\Gamma_{\mathcal{O} \cup \{i\}}((1 - \kappa)\nu))}{\rho_i(\Gamma_{\mathcal{P}}(\kappa\nu))} & \text{if } i \in \mathcal{P}. \end{cases} \quad (4)$$

2.3 Competitive Equilibrium

Notice that a CP's joining decision to a service class might increase the congestion level and reduce the throughput of flows of that service class; however, if the number of CPs in a service class is big, an additional CP i 's effect will be marginal. Analogous to the *pricing-taking assumption* [9] in a competitive market, we can make a *throughput-taking assumption* for the CPs as follows.

Assumption 4 (CONGESTION TAKING). *In a competitive equilibrium, for any service class \mathcal{X} , any CP $i \notin \mathcal{X}$ uses $\Gamma_{\mathcal{X}}$ as an estimate of the ex-post congestion $\Gamma_{\mathcal{X} \cup \{i\}}$ it will face in the decision-making.*

Based on the above congestion-taking assumption, we can define a competitive equilibrium of the CPs as follows.

Definition 3. *A strategy profile $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$ is a competitive equilibrium of a game $(M, \mu, \mathcal{N}, s_I)$, if*

$$\frac{v_i - c}{v_i} \begin{cases} \leq \frac{\rho_i(\Gamma_{\mathcal{O}}((1 - \kappa)\nu))}{\rho_i(\Gamma_{\mathcal{P}}(\kappa\nu))} & \text{if } i \in \mathcal{O}, \\ > \frac{\rho_i(\Gamma_{\mathcal{O}}((1 - \kappa)\nu))}{\rho_i(\Gamma_{\mathcal{P}}(\kappa\nu))} & \text{if } i \in \mathcal{P}. \end{cases} \quad (5)$$

Theorem 3. *If $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$ is an equilibrium of a game $(M, \mu, \mathcal{N}, s_I)$, it is also a same type of equilibrium (Nash or competitive) of a game $(\xi M, \xi \mu, \mathcal{N}, s_I)$ for any $\xi > 0$.*

Proof of Theorem 3: By Corollary 1, $(\mathcal{O}, \mathcal{P})$ would induce the same level of congestion $\Gamma_{\mathcal{N}}$ in a linearly scaled system $(\xi M, \xi \mu, \mathcal{N}, s_I)$. Therefore, the right hand sides of both 4 and 5 do not change. Since the left hand sides of both 4 and 5 remain the same, all the equilibrium conditions are satisfied. ■

2.4 Finding a Competitive Equilibrium

We will focus on competitive equilibria rather than Nash equilibria for two reasons. First, because the number of CPs in practice is big, the congestion-taking assumption is valid. Second, the *common knowledge* assumption [9] for reaching Nash equilibria might be questionable, because CPs rarely know the characteristics of all other CPs in practice. Nevertheless, many results apply for both equilibrium definitions.

The difficulty of finding a competitive equilibrium comes from the large scale of number of CPs. Through exhaustive search, we need to evaluation $2^{|\mathcal{N}|}$ possible scenarios. Here, we first introduce a condition under which we can efficiently obtain a competitive equilibrium and then we explore a more general method to find a competitive equilibrium.

2.4.1 A special form of $\rho_i(\Gamma)$

Assumption 5. *For any two levels of congestion Γ_1 and Γ_2 , any congestion function $\rho_i(\Gamma)$ satisfies*

$$\frac{\rho_i(\Gamma_1)}{\rho_i(\Gamma_2)} = F_i(G(\Gamma_1, \Gamma_2)),$$

where $F_i(\cdot)$ is continuous and invertible, and $G(\Gamma_1, \Gamma_2)$ is a common continuous function among all CPs that is non-increasing in Γ_1 and non-decreasing in Γ_2 .

Two examples of ρ_i that satisfy the above assumption are

$$\rho_i(\Gamma) = \hat{\theta}_i e^{-(\beta_i)\Gamma}, \quad \text{and} \quad \rho_i(\Gamma) = \hat{\theta}_i (\beta_i)\Gamma,$$

where β_i can serve as a differentiating parameter that reflects CP i 's to sensitivity congestion. If the above two forms of

ρ_i are adopted, the corresponding demand functions $d_i(\Gamma)$ under the proportional rate mechanism would be

$$d_i(\Gamma) = (\Gamma + 1)e^{-(\beta_i)\Gamma}, \quad \text{and} \quad d_i(\Gamma) = (\Gamma + 1)(\beta_i)\Gamma.$$

Similarly, the corresponding demand functions $d_i(\Gamma)$ under a max-min end-to-end congestion mechanism would be

$$d_i(\Gamma) = \max\{1, \hat{\theta}_i \Gamma\} e^{-(\beta_i)\Gamma}, \quad \text{and} \quad d_i(\Gamma) = \max\{1, \hat{\theta}_i \Gamma\} (\beta_i)\Gamma.$$

If Assumption 5 is satisfied, we denote ζ_i as a relative priority of CP i under a given charge c , defined by

$$\zeta_i = \begin{cases} F_i^{-1}\left(\frac{v_i - c}{v_i}\right) & \text{if } v_i > c, \\ +\infty & \text{otherwise.} \end{cases} \quad (6)$$

The smaller priority value a CP has, the larger chance it will end up in the premium class in an equilibrium. We define $\mathcal{R}_i^+ = \{i' : \zeta_{i'} > \zeta_i\}$ and $\mathcal{R}_i^- = \{i' : \zeta_{i'} < \zeta_i\}$ as the set of CPs that have bigger and smaller priorities than that of CP i .

Lemma 2. *If $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$ is a competitive equilibrium of a game $(M, \mu, \mathcal{N}, s_I)$ with $s_I = (\kappa, c)$, then for any $i \in \mathcal{O}$ and $j \in \mathcal{P}$, we have $\zeta_i > \zeta_j$, $\mathcal{R}_i^+ \subset \mathcal{O}$ and $\mathcal{R}_j^- \subset \mathcal{P}$.*

Lemma 2 reveals the structure of a competitive equilibrium under which \mathcal{O} always contains CPs with higher values of ζ_i . Therefore, to find a competitive equilibrium, we can first sort the values of ζ_i in a descending order and relabel the CPs such that $\zeta_i \geq \zeta_j$ if $i < j$, and then, find a partition of the sorted CPs under which the CPs with lower/higher indices will be in \mathcal{O}/\mathcal{P} . We denote \mathcal{H}_L and \mathcal{T}_L as the set of the first L and the last L CPs in the sorted list respectively. If we make a partition L such that the first L CPs are put in the ordinary class and the last $N - L$ CPs are put in the premium class, the congestion difference of the two classes is

$$\Delta_L = \Gamma_{\mathcal{H}_L} - \Gamma_{\mathcal{T}_{N-L}}. \quad (7)$$

By Corollary 1, we know that the congestion level Γ_S is non decreasing with the size of the set S of CPs. Therefore, $\{\Delta_0, \dots, \Delta_N\}$ is a non decreasing sequence.

Lemma 3. *Suppose CPs are sorted by ζ_i in a descending order. If all ζ_i s have distinct values except for $+\infty$, then there exist at most one CP $l \in \mathcal{N}$ such that $\Delta_{l-1} < \zeta_l < \Delta_l$. If no such CP l exists, the game $(M, \mu, \mathcal{N}, s_I)$ has a unique competitive equilibrium; otherwise, it does not have any.*

The condition of distinct ζ_i s excludes multiple equilibria. If there exist $\zeta_i = \zeta_j < +\infty$, then i and j might switch places to generate symmetric competitive equilibria. Without loss generality, we can relabel the CPs such that CPs in \mathcal{O} will have smaller indices and still keep ζ_i in a descending order. Because ζ_i is a non-increasing sequence and Δ_L is a non decreasing sequence, there exists at most one crossing point where a CP l might find itself better off in \mathcal{P} when it is in \mathcal{O} and vice versa. This happens because when a CP moves from one service class to another, the congestion differences Δ_L make jumps. Asymptotically, when N increases, the

sequence of Δ_{L_S} will be smoother and a unique competitive equilibrium will almost always exist (like the competitive market equilibrium under the price-taking assumption in classic economics). To resolve this technical problem, we assume that if such a CP l exists, it will eventually calculate its real ex-post congestion levels as in the Nash equilibrium condition (4) and stabilize itself in one service class. We implement the following binary-search algorithm to find the unique equilibrium for our numerical analysis purposes.

Finding CPs' Competitive Equilibrium (M, μ, \mathcal{N}, s_I)

1. Sort CPs based on ζ_i in a descending order;
 2. $L_1=0; L_2 = N; L = \frac{1}{2}(L_1 + L_2)$;
 3. **while** $L_2 > L_1$
 4. **if** $\Gamma_{\mathcal{H}_L} - \Gamma_{\mathcal{T}_{N-L+1}} < \zeta_L$ **then** $L_2 = L - 1$;
 5. **if** $\Gamma_{\mathcal{H}_{L+1}} - \Gamma_{\mathcal{T}_{N-L}} > \zeta_{L+1}$ **then** $L_1 = L$;
 6. **if** $L == \frac{1}{2}(L_1 + L_2)$ **then** $L_1 = L_2 = L$;
 7. **else** $L = \frac{1}{2}(L_1 + L_2)$;
 8. $\mathcal{O} = \mathcal{H}_L; \mathcal{P} = \mathcal{T}_{N-L}$;
 9. **return** (\mathcal{O}, \mathcal{P});
-

2.4.2 General forms of $\rho_i(\Gamma)$

For general forms of $\rho_i(\Gamma)$ s, we cannot sort the CPs on a one-dimensional space to search for a competitive equilibrium. However, we could still search in a two-dimensional "congestion space" as follows.

Finding CPs' Competitive Equilibrium (M, μ, \mathcal{N}, s_I)

1. Initialize $\Gamma[0] = (\Gamma_{\mathcal{O}}^{[0]}, \Gamma_{\mathcal{P}}^{[0]})$
 2. Calculate induced equilibrium ($\mathcal{O}_{[0]}, \mathcal{P}_{[0]}$) given $\Gamma[0]$;
 3. $t = 0$;
 4. **do**
 5. $\Gamma'[t] = (\Gamma_{\mathcal{O}_{[t]}}', \Gamma_{\mathcal{P}_{[t]}}')$;
 6. $\Gamma[t+1] = \Gamma[t] + g[t](\Gamma'[t] - \Gamma[t])$;
 7. $t = t + 1$;
 8. Calculate ($\mathcal{O}_{[t]}, \mathcal{P}_{[t]}$) based on $\Gamma[t]$;
 9. **until** $t > T$ or ($\mathcal{O}_{[t]}, \mathcal{P}_{[t]} == (\mathcal{O}_{[t-1]}, \mathcal{P}_{[t-1]})$);
 10. **return** ($\mathcal{O}_{[t]}, \mathcal{P}_{[t]}$);
-

The algorithm starts with an estimation $\Gamma[0] = (\Gamma_{\mathcal{O}}^{[0]}, \Gamma_{\mathcal{P}}^{[0]})$ on the congestion levels in both service classes. At each time step t , based on the congestion-taking assumption and the congestion estimation $\Gamma[t]$, the algorithm calculate an induced competitive equilibrium ($\mathcal{O}_{[t]}, \mathcal{P}_{[t]}$) (line 2 and 8) and the corresponding real congestion level $\Gamma'[t] = (\Gamma_{\mathcal{O}_{[t]}}', \Gamma_{\mathcal{P}_{[t]}}')$ (line 5). The algorithm updates the congestion $\Gamma[t+1]$ based on the previous estimate $\Gamma[t]$ and the induced real congestion $\Gamma'[t]$ (line 6). Notice that the algorithm is flexible that a step size $g[t]$ for each round (line 6) and a maximum number of rounds T (line 9) can be specified for the algorithm to trade-off between convergence time and accuracy.

3. RELATED WORK AND CONCLUSIONS

From the game model for service differentiation developed in the last section, we can further analyze the strategic decision of the ISP and the resulting welfare for different players. In [8], we focus on rate allocation mechanisms of the system and further analyze the oligopolistic competition among the ISPs. The results reveal some insight about the appropriateness of imposing network neutrality regulations.

Besides the use of analyzing differentiated service classes, the congestion equilibrium is more fundamental, upon which we can build more sophisticated higher level interaction models. Moreover, the notion of congestion can be thought as a generalization of price which determines an equilibrium for competitive economy of commodity markets [9].

Acknowledgments: This work was partially supported by the HSSP grant from Advanced Digital Sciences Center, MOE AcRF grant R-252-000-448-133, the National Science Foundation grants CNS-1017934 and CCF-1139915. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

4. REFERENCES

- [1] Google corporate website. <http://www.google.com/corporate/tech.html>.
- [2] Netflix technology blog. <http://techblog.netflix.com/2011/01/netflix-performance-on-top-isp-networks.html>.
- [3] M. Campbell and J. Browning. Apple, Google asked to pay up as mobile operators face data flood. *Bloomberg News*, December 7 2010.
- [4] D. M. Chiu and R. Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN Systems*, 17(1).
- [5] D. Clark. The design philosophy of the DARPA Internet protocols. *ACM SIGCOMM Computer Communication Review*, 18(4), August 1988.
- [6] C. Courcoubetis and R. Weber. *Pricing Communication Networks: Economics, Technology and Modelling*. John Wiley & Sons Ltd., 2003.
- [7] 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.
- [8] R. T. B. Ma and V. Misra. The Public Option: a non-regulatory alternative to network neutrality. *Proceedings of the 7th International Conference on emerging Networking Experiments and Technologies (CoNEXT)*, 2011.
- [9] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- [10] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5), October 2000.
- [11] A. Odlyzko. Paris metro pricing for the Internet. *Proceedings of ACM Conference on Electronic Commerce*, 1999.
- [12] M. J. Osborne and A. Rubinstein. *A course in game theory*. The MIT Press Cambridge, Massachusetts, 1994.
- [13] N. Shetty, G. Schwartz, and J. Walrand. Internet QoS and regulations. *IEEE/ACM Transactions on Networking*, 18(6), December 2010.