Internet Transport Economics: Model and Analysis

Richard T. B. Ma
Department of Computer Science
National University of Singapore
tbma@comp.nus.edu.sg

ABSTRACT

With the rise of video streaming and cloud services, the Internet has evolved into a content-centric service network; however, quality of service (QoS) is still a major concern for the content providers. Because quality degradation is influenced by 1) the capacities of links along the routes used for content delivery and 2) the amount of competing traffic across these links, it is very difficult to diagnose.

In this paper, we establish a novel model to study how business decisions such as capacity planning, routing strategies and peering agreements affect QoS in terms of the end-to-end delays and drop rates of Internet routes. In particular, we take an economics perspective of the Internet transport service and model its supply of network capacities and demands of throughput driven by network protocols. We show that a macroscopic network equilibrium always exists and its uniqueness can be guaranteed under various scenarios. We analyze the impacts of user demands and resource capacities on the network equilibrium and provide implications of Netflix-Comcast type of peering on the QoS of users. We demonstrate that our framework can be used as a building block to understand the routing strategies under a Wardrop equilibrium and to enable further studies such as Internet peering and in-network caching.

CCS CONCEPTS

• **Networks** → **Network economics**; **Network performance modeling**; **Network performance analysis**.

ACM Reference Format:
Richard T. B. Ma. 2020. Internet Transport Economics: Model and Analysis. In *The Twenty-first ACM International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (Mobihoc '20), October 11–14, 2020, Boston, MA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3397166.3409137

1 INTRODUCTION

The Internet is the underpinning of today's digital economy. As a packet switching network, it provides transport/delivery services for data packets, through which innovative applications and online services across the globle are enabled. For example, Netflix [10] streams videos to residential homes over the Internet and now accounts for up to 34% of peak U.S. downstream traffic. Such content services require low delay and high throughput and are quite sensitive to quality of service (QoS); however, the causes and effects of quality degradation are often difficult to diagnose and predict if occurs. In January 2014, the download speed of Netflix users behind the access provider Comcast was found to be dropped by 25%. Netflix blamed Comcast [7] for throttling traffic, but Comcast claimed that Netflix was sending high rates of traffic to intentionally congest the peering links. Although this peering dispute was resolved by Netflix paying Comcast to reach a premium peering agreement, this paid prioritization practice had raised concerns about net neutrality [32], whose impacts on the Internet and its evolution are largely unknown. Consequently, the U.S. Federal Communications Commission (FCC) decided to closely monitor but exempt these non-neutral practices from its recent Open Internet Order [1], because it lacked background "in the Internet traffic exchange context." Both the debate of net neutrality and peering disputes are issues of network economics, which boil down to understanding the policy implications on the QoS for Internet applications.

However, it is challenging to faithfully characterize the QoS of an application, because it is influenced by the amount of competing traffic across the links, and their link capacities, along the routes used for content delivery, which are not determined by a single administrative domain but collectively controlled by thousands of interconnected autonomous systems (ASes). Therefore, a comprehensive QoS model has to capture the complex business decisions of ASes, e.g., the capacity planning of the Internet service providers (ISPs), the routing strategies of the content providers (CPs) and the peering agreements between CPs and ISPs.

In this paper, we take an economics approach to study the Internet transport service and its QoS. *We develop a holistic model with an analytical framework that enables the analyses of network economics such as the impacts of inter-dependent business decisions of ASes on the resulting QoS.* Unlike traditional transport economics [29] that studies the movement of people and goods over space and time, Internet transport economics studies the movement of streams of data packets that creates information services. In particular, we consider quality metrics of the end-to-end delay and the drop rate of such services and model their supply and demand based on the characteristics of network resources deployed by ASes and network protocols used by applications, respectively. We formulate and study the network equilibrium under which the quality metrics are determined in a steady state. Our framework and fixed-point analysis generalize prior work [9, 13] that model the average rate of TCP flows, where neither the existence nor the uniqueness of solution has been established. Our analytical contributions include the following.

- We prove the existence of a general network equilibrium and derive the conditions for its uniqueness (Theorem 1).

Mobihoc '20, October 11–14, 2020, Boston, MA, USA

ACM ISBN 978-1-4503-8015-7/20/10…$15.00
https://doi.org/10.1145/3397166.3409137

- We prove the uniqueness of network equilibrium for any linear topology or lossless network (Theorem 2 and 4).
- We analyze the impact of user population and network capacity on the network equilibrium (Theorem 3 and 5).
- We demonstrate that the network equilibrium can be used as a building block to understand the routing strategies of ASes under a Wardrop equilibrium (Theorem 6).

Our analytical results provide new implications of the peering relationship between Netflix and Comcast on users' QoS as follows.

- Even if Comcast dedicate isolated resources for Netflix, the QoS will degrade as Netflix attracts more users.
- Under any fixed number of users, if Netflix increases capacities, it will increase its aggregate throughput but make downstream resources at Comcast more congested.
- Under a general lossless network, the increase of any resource capacity might not always improve the QoS on any route due to the inter-dependency among the routes.

We foresee that our framework can be used as a fundamental physical model for the Internet transport ecosystem, on top of which strategic routing, peering and pricing decisions of the ASes and their impacts on QoS can be further studied.

Notation	Semantics
$\mathcal{K}, k$	Set of network resources and a resource $k \in \mathcal{K}$
d_k, g_k	Delay and gain of network resource k
Ψ_k, Ψ_k^g	Delay and gain functions of resource k
λ_k, ϕ_k	Generic and aggregate input rate to resource k
μ_k, τ_k	Capacity and throughput of resource k
$\boldsymbol{d}, \boldsymbol{g}, \boldsymbol{\phi}$	Vectors of delay, gain and input rate of resources
$\mathcal{L}, l$	Set of network routes and a route $l \in \mathcal{L}$
D_l, G_l	End-to-end delay and gain of network route l
φ_l	Aggregate input rate to network route l
Φ_l	Aggregate input rate function of route l
$\boldsymbol{\varphi}, \boldsymbol{D}, \boldsymbol{G}$	Vectors of delay, gain and input rate of routes
$\mathcal{I}, i$	Set of content providers (CPs) and a CP $i \in \mathcal{I}$
s_{il}	Number of users of CP i served via route l
Λ_i	Average per-user sending rate function of CP i

Table 1: Summary of notation used in this paper.

2 MODELS OF SUPPLY AND DEMAND

The Internet interconnects billions of hosts or end-systems by a network of communication links and packet switches. To study its fundamental economics, we first need to understand the commodities produced and consumed from the Internet. General economic commodities comprise goods and services. Instead of producing physical goods, the Internet provides transport services for information goods, whose characteristics are defined by performance metrics such as throughput, delay and drop rate.

In this section, we start modeling the physical behaviors and mechanisms of network resources and protocols that fundamentally drive the supply and demand of the Internet transport services.

2.1 Supply-Side Model of Network Resources

The transportation of packets across the Internet is enabled by switching devices, e.g., routers and switches, and links, e.g., fiber optics. Despite the differences in technical characteristics, we conceptualize *network resource* to be any physical entity that enables the transportation of data packets from one physical location to another. Depending on the modeling granularity, a network resource may represent a physical link, an AS or an ISP in practice.

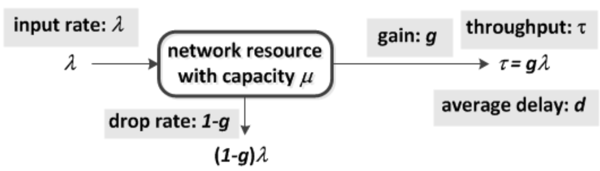

Figure 1: Illustration of a generic network resource.

Figure 1 illustrates the basic behavior of a network resource. Given a capacity μ and an input rate λ, a fraction of packets might be dropped due to congestion and buffer overflow. We denote the *gain* of transmission by g and the drop rate of packets by $1-g$; and therefore, define the rate of successful transmission or *throughput* by $\tau = g\lambda$. Each transmitted packet spends some time going through the network resource. We denote this average sojourn time or *delay* of packets by d.

We consider a generic network that comprises a set $\mathcal{K}$ of network resources. Since both delay and gain depend on the capacity of and the input rate to a network resource, we assume that the delay d_k and gain g_k of each resource $k \in \mathcal{K}$ are functions of its capacity μ_k and input rate λ_k as follows.

Assumption 1 (Delay and Gain Monotonicity). *The delay and gain of any resource $k \in \mathcal{K}$ can be expressed as the functions*

$$d_k = \Psi_k(\lambda_k, \mu_k) \quad and \quad g_k = \Psi_k^g(\lambda_k, \mu_k), \tag{1}$$

where $\Psi_k(\lambda_k, \mu_k)$ is non-decreasing in λ_k and non-increasing in μ_k; and $\Psi_k^g(\lambda_k, \mu_k)$ is non-increasing in λ_k and non-decreasing in μ_k. Both $\Psi_k(\lambda_k, \mu_k)$ and $\Psi_k^g(\lambda_k, \mu_k)$ are differentiable.

Assumption 1 states that when the capacities expand or the input rates reduce, network resources become less congested and thus provide lower delays and higher gains, and vice-versa.

Since network elements often consist of buffers with queues, e.g., in the output ports of routers, queueing models can be used to characterize the performance of network resources. We show a couple of examples of such queueing models in the following.

The $M/G/1$ model: Under a general service time of packets, the delay is specified by the Pollaczek-Khinchine formula as

$$d_k = \Psi_k(\lambda_k, \mu_k) = \frac{1}{\mu_k} + \frac{\mathbb{E}[S^2]}{2} \frac{\lambda_k}{1 - \lambda_k/\mu_k}, \quad \forall \lambda_k < \mu_k.$$

Notice that the capacity μ_k determines the mean service time $\mathbb{E}[S] = \mu_k^{-1}$, which is the lower-bound for delay. Furthermore, under any admissible input rate λ_k, the delay d_k depends on the second moment of service time $\mathbb{E}[S^2]$. Thus, the variability of service time under the queueing model can be used to model network resources that have different delay characteristics.

Given a fixed capacity μ_k and a desirable delay d_k, we can also derive the maximum amount of admissible rate as

$$\lambda_k^{max} = \frac{\left(d_k - \mu_k^{-1}\right)}{\mathbb{E}[S^2]/2 + \left(d_k - \mu_k^{-1}\right)\mu_k^{-1}}, \quad \forall d_k > \mu_k^{-1},$$

where a lower variability of $\mathbb{E}[S^2]$ implies that the network resource can accommodate higher rates under delay constraints.

One limitation of the $M/G/1$ model is that it assumes an infinite queue and cannot characterize the delays and losses under heavily loaded scenarios. A natural extension is the $M/G/1/K$ model that captures the buffer size or the applied active queue management mechanisms. Next, we illustrate an example where the service time follows a gamma distribution.

The $M/\Gamma/1/K$ model: This queueing model assumes that the system can accommodate at most K packets at any time and the service time distribution is governed by a shape parameter α that determines its mean and second moment as

$$\mathbb{E}[S] = \mu^{-1} \quad \text{and} \quad \mathbb{E}[S^2] = \left(\alpha^{-1} + 1\right)\mu^{-2}.$$

To better illustrate the behaviors of delay and gain of a network resource k under the $M/\Gamma/1/K$ model, we normalize the capacity to be $\mu_k = 1$ without loss of generality.

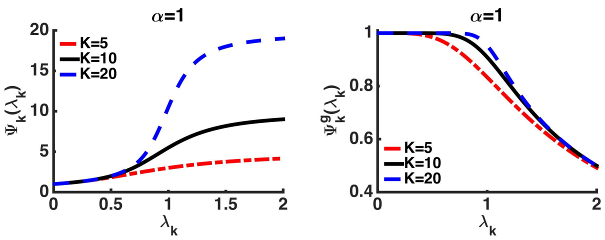

Figure 2: Delay and gain of a network resource under varying input rates.

Figure 2 plots the delay d_k and gain g_k as functions of the input rate λ_k in the left and right sub-figures, respectively. The shape parameter α is fixed to be 1, while $K = 5, 10$ and 20 in the three curves. In general, we observe that the delay increases but the gain decreases with the input rate, satisfying Assumption 1. In particular, when K becomes larger, the network resource achieves a higher gain, as well as a higher throughout, at a cost of higher delay. This reflects the general tradeoff between delay and throughput when active queue management (AQM) mechanisms [8] are employed to limit the length of queue.

Similar to Figure 2, Figure 3 plots the delay d_k and gain g_k under $K = 10$, while $\alpha = 0.1, 1$ and 10 in the three curves. We observe that when α increases, the network resource achieves a higher gain; however, the delay increases with α only under a heavy traffic region where $\lambda_k > \mu_k$. The parameter α can be used to model various elasticity characteristics of network resources, i.e., how the

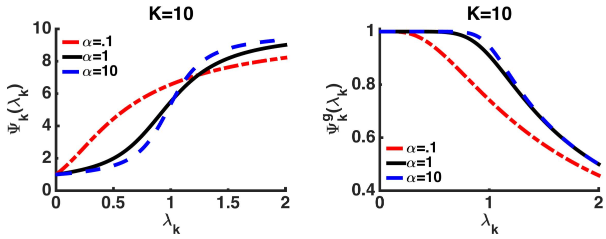

Figure 3: Delay and gain of a network resource under varying input rates.

throughput and gain of a resource respond to the changes in the external input rate.

Based on Assumption 1, the supply of service provided by any resource k can be characterized by the quality metrics d_k and g_k. However, both are not exogenous, but endogenously determined by the quantity metric λ_k. Because $\tau_k = g_k\lambda_k$, the realized service can be characterized by a pair (d_k, τ_k), capturing both the quality and quantity of the service.

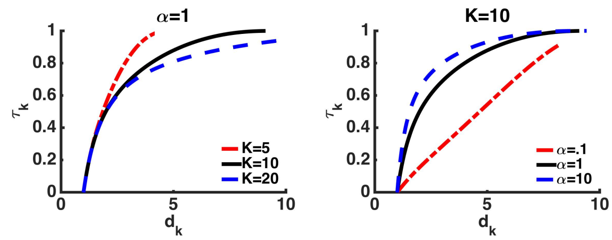

Figure 4: Maximum achievable throughput of under varying desirable delay.

Figure 4 plots the maximum throughput τ_k resource k can accommodate under a desirable delay d_k that varies along the x-axis. We fix α and vary K in the left sub-figure and fix K and vary α in the right sub-figure. Although a larger K reduces the drop rate in general, we observe in the left sub-figure that it increases the delay when accommodating the same amount of throughput and accommodates less throughput under the same desirable delay. In the right sub-figure, we observe that a larger α reduces the delay when accommodating the same amount of throughput and accommodates more throughput under the same delay. Notice that each curve shows the characteristics of the service provided by the resource, since each point (d_k, τ_k) represents a possible resulting service under certain input rate.

Before we close this subsection, we would like to emphasize that neither is the queueing model the only means to describe network resources, nor is it applied in a conventional manner. For example, the service time distribution in a queueing model is used under our context to capture the characteristics of a network resource in terms of its delay and gain in response to the input rates.

2.2 Demand-Side Model of Users and Application Protocols

The Internet can be conceptually viewed as a two-sided market [25] that connects end-users to online content providers. Becuase the Internet transport services are based on the end-to-end routes that consist of multiple network resources, we define a route l to be a non-empty set $l \subseteq \mathcal{K}$ of resources that are serially connected. Based on the additivity and multiplicative properties of delay and gain, we define the delay D_l and gain G_l of a route l as

$$D_l = \sum_{k \in l} d_k \quad \text{and} \quad G_l = \prod_{k \in l} g_k. \tag{2}$$

We denote the set of content providers (CPs) by $\mathcal{I}$ and the average sending rate of any CP $i \in \mathcal{I}$ to an active user by λ_i, which can be regarded as the user's demand for transport services. This demand is driven by end-users running various application and network protocols and is influenced by the service quality in terms of delay and gain of the routing paths. For example, the PFTK formula [22] characterizes the rate of a single TCP Reno flow as a function of delay and gain as

$$\lambda_i(D, G) \approx (2D)^{-1}\sqrt{\frac{3}{4(1-G)}} + o\left(\sqrt{\frac{1}{1-G}}\right), \tag{3}$$

which states that the flow rate is inversely proportional to the round-trip time, measured by $2D$, and the square root of drop rate, measured by $\sqrt{1-G}$. Under serve network congestion that induces high delay and drop rate, active users might become impatience and stop using the transport service. We define $\eta_i(D, G)$ to be the probability that an active user will still use the service under the delay D and gain G. For any CP $i \in \mathcal{I}$, we define the average per-user sending rate by

$$\Lambda_i(D, G) = \eta_i(D, G)\lambda_i(D, G). \tag{4}$$

Assumption 2 (Demand Rate Monotonicity). *The average per-user sending rate $\Lambda_i(D, G)$ of any CP $i \in \mathcal{I}$ is a differentiable function, decreasing in D and non-decreasing in G.*

Assumption 2 states that the demand for transport service will not increase if the delay or gain of the route deteriorates.

The iso-elastic demand model: Although Assumption 2 is very general, one neat model for the demand function can be

$$\Lambda_i(D, G) \propto D^{-\beta}(1-G)^{-\gamma},$$

where the parameters β and γ define the elasticity of demand rate Λ_i with respective to delay D and gain G, which capture the demand rate's sensitivity to the round-trip time and drop-rate, respectively. The PFTK formula (3) has shown that the elasticity parameters satisfy $\beta = 1$ and $\gamma = 1/2$ for an active TCP Reno flow. For network protocols that are less sensitive to delay or loss, e.g., UDP, the values of the corresponding parameters will be lower[1]. It is interesting to notice that from an economics point of view, an *elastic demand function* corresponds to high values of β and γ; however from a networking perspective, it models the *inelastic traffic* [28], e.g., video streaming, that is intolerable to delay and packet losses.

[1]If the protocol does not respond to a metric, its parameter equals zero.

3 NETWORK EQUILIBRIUM

In this section, we study the equilibria of a network system, under which both the supply and demand of transport services, discussed in the previous section, are balanced in steady states.

We denote the set of all feasible routes in the system as $\mathcal{L}$, which are determined by the collective peering decisions of ISPs with other ISPs and CPs and materialized via the BGP inter-domain routing protocol. Any route $l \in \mathcal{L}$ consists of an ordered sequence of network resources that lead to a number of end-users in a geographical region. For any CP $i \in \mathcal{I}$, we denote the number of end-users it serves via route $l \in \mathcal{L}$ by s_{il}. If a route l is not available to CP i, possibly due to its peering relationships with ISPs, or CP i chooses not to use route l due to performance reasons, s_{il} is set to be zero without loss of generality. As a result, the aggregate sending rate from all CPs along route l to end-users can be defined as

$$\varphi_l = \Phi_l(D_l, G_l) = \sum_{i \in \mathcal{I}} s_{il}\Lambda_i(D_l, G_l), \tag{5}$$

which is a function of route l's service quality in terms of delay D_l and gain G_l. Given the rate φ_l of any route $l \in \mathcal{L}$, the input rates to any resource $k \in \mathcal{K}$ can be derived by

$$\phi_k = \sum_{l \ni k} \varphi_l \prod_{\kappa \in l(k)} g_\kappa, \tag{6}$$

where $l(k)$ defines the set of resources on route l before k. Equation (6) states that the aggregate input rate to any resource k equals the sum over the sending rates φ_l whose route utilizes resource k, i.e., $k \in l$, discounted by the gains of the resources $l(k)$ preceeding k along various routes. In particular, if $l(k)$ is empty, i.e., k is the first resource on route l, the nullary product equals the multiplicative identity 1 by convention.

By far a network system can be defined by a triple $(\mathcal{I}, \mathcal{K}, \mathcal{L})$ that describe the sets of CPs, network resources and routes. Under a steady state, the rates ϕ_k and φ_l over the resources and routes determine and are influenced by the quality metrics (d_k, g_k) and (D_l, G_l) at the resource- and route-level, respectively. We denote $\boldsymbol{\phi}$, $\boldsymbol{d}$, and $\boldsymbol{g}$ as the vectors of resource-level metrics and $\boldsymbol{\varphi}$, $\boldsymbol{D}$ and $\boldsymbol{G}$ as the corresponding vectors at the route level. By using Equations (1), (2), (5) and (6), we can define an equilibrium of a network system $(\mathcal{I}, \mathcal{K}, \mathcal{L})$ as follows.

Definition 1 (Network Equilibrium). *For any system $(\mathcal{I}, \mathcal{K}, \mathcal{L})$, a tuple $(\boldsymbol{\phi}, \boldsymbol{d}, \boldsymbol{g}, \boldsymbol{\varphi}, \boldsymbol{D}, \boldsymbol{G})$ is an equilibrium if and only if*

$$\begin{cases} \varphi_l = \Phi_l(D_l, G_l) = \sum_{i \in \mathcal{I}} s_{il}\Lambda_i(D_l, G_l), & \forall\, l \in \mathcal{L}. \\ (d_k, g_k) = \left(\Psi_k(\phi_k, \mu_k), \Psi^g_k(\phi_k, \mu_k)\right), & \forall\, k \in \mathcal{K}. \\ (D_l, G_l) = \left(\sum_{k \in l} d_k, \prod_{k \in l} g_k\right), & \forall\, l \in \mathcal{L}. \\ \phi_k = \sum_{l \ni k} \varphi_l \prod_{\kappa \in l(k)} g_\kappa, & \forall\, k \in \mathcal{K}. \end{cases} \tag{7}$$

The first two equations of Definition 1 specify the demand of transport services at the route level and the corresponding supply at the resource level, respectively. The last two equations specify

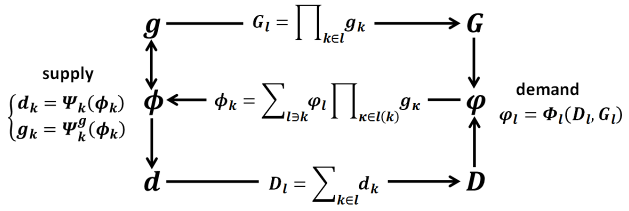

Figure 5: An illustration of the network equilibrium.

the physics of delay, gain and aggregate input rate that links the metrics between the resource and route levels.

Figure 5 visualizes the complex relationship among different metrics under an equilibrium. The right shows that the demand $\boldsymbol{\varphi}$ is driven by the route level metrics $\boldsymbol{D}$ and $\boldsymbol{G}$; the left shows that the endogenously supplied quality metrics $\boldsymbol{d}$ and $\boldsymbol{g}$ are driven by the resource-level demand $\boldsymbol{\phi}$. In between the demand and supply, relationships of the gains, rates and delays at the resource and route levels are shown from top to bottom.

Notice that if the vector $\boldsymbol{\phi}$ under an equilibrium is known, all others can be uniquely determined. To compactly express equilibria in a vector form and characterize them, we define a vector $\Phi(\boldsymbol{\phi})$ of aggregate sending rates along all the routes as a function of $\boldsymbol{\phi}$, where each entry $\Phi_l(\boldsymbol{\phi})$ is defined as

$$\Phi_l(\boldsymbol{\phi}) = \sum_{i \in \mathcal{I}} s_{il} \Lambda_i \left(\sum_{k \in l} \Psi_k(\phi_k, \mu_k), \prod_{k \in l} \Psi^g_k(\phi_k, \mu_k) \right). \quad (8)$$

We also define an $|\mathcal{L}| \times |\mathcal{K}|$ matrix $H(\boldsymbol{\phi})$ as functions of $\boldsymbol{\phi}$, where each entry $H_{lk}(\boldsymbol{\phi})$ is the effective gain along route l before arriving resource k under $\boldsymbol{\phi}$, defined as

$$H_{lk}(\boldsymbol{\phi}) = \mathbb{1}_{\{k \in l\}} \prod_{\kappa \in l(k)} \Psi^g_\kappa(\phi_\kappa, \mu_\kappa). \quad (9)$$

The following result characterizes the existence and uniqueness of equilibrium based on the properties of $\Phi(\cdot)$ and $H(\cdot)$.

Theorem 1. *Under Assumption 1 and 2, for any system* $(\mathcal{I}, \mathcal{K}, \mathcal{L})$, *there always exists an equilibrium that satisfies*

$$\boldsymbol{\phi} = H(\boldsymbol{\phi})^T \Phi(\boldsymbol{\phi}). \quad (10)$$

Furthermore, let $\phi^{max}_k = \sum_{l \ni k} \Phi_l(0, 1)$ *for all* $k \in \mathcal{K}$. *The equilibrium is unique, if for all* $\boldsymbol{\phi} \in \times_{k \in \mathcal{K}} [0, \phi^{max}_k]$

$$\left| \nabla_{\boldsymbol{\phi}} \Phi(\boldsymbol{\phi}) H(\boldsymbol{\phi}) + \nabla_{\boldsymbol{\phi}} H(\boldsymbol{\phi}) \Phi(\boldsymbol{\phi}) - I \right| \neq 0, \quad (11)$$

where ∇ *and* I *are the gradient operator and identity matrix.*

Proof of Theorem 1: $\Phi(\boldsymbol{\phi})$ is a vector of functions $\varphi_l(\boldsymbol{\phi})$, each of which is a composite of the first three equations in Definition 1. Therefore, the fourth equation can be written in a vector form as Equation (10) to define an equilibrium.

Because the maximum input rate to resource k is ϕ^{max}_k, any equilibrium $\boldsymbol{\phi}$ lies in the feasible domain $\times_{k \in \mathcal{K}} [0, \phi^{max}_k]$. To show the existence of equilibrium, let $F(\boldsymbol{\phi}) = H(\boldsymbol{\phi})^T \Phi(\boldsymbol{\phi})$. $F(\cdot)$ is a continuous mapping from the convex compact subset $\times_{k \in \mathcal{K}} [0, \phi^{max}_k]$ of a Euclidean space to itself; and therefore, by Brouwer fixed-point theorem [16], there always exists a fixed point that satisfies $\phi = F(\phi) = H(\boldsymbol{\phi})^T \Phi(\boldsymbol{\phi})$.

For any active resource k, $d_k > 0$ and the maximum rate ϕ^{max}_k cannot be achieved under an equilibrium, because any route $l \ni k$ has a positive delay D_l and by Assumption 2, its demand $\Lambda_i(D_l, G_l)$ will be lower than $\Lambda_i(0, G_l)$, which is lower than its contribution $\Lambda_i(0, 1)$ to ϕ^{max}_k. As a result, any boudary value of $\times_{k \in \mathcal{K}} [0, \phi^{max}_k]$ cannot be an equilibrium. By the chain rule, the gradient of $F(\boldsymbol{\phi})$ can be written as $\nabla_{\boldsymbol{\phi}} F(\boldsymbol{\phi}) = \nabla_{\boldsymbol{\phi}} \Phi(\boldsymbol{\phi}) H(\boldsymbol{\phi}) + \nabla_{\boldsymbol{\phi}} H(\boldsymbol{\phi}) \Phi(\boldsymbol{\phi})$ and Equation (11) guarantees that $\mathbf{1}$ cannot be an eigenvalue of $\nabla_{\boldsymbol{\phi}} F(\boldsymbol{\phi})$. Finally, by using the Kellogg's fixed-point theorem [15], we conclude that the uniquess of equilibrium can be guaranteed by Equation (11) over the feasible domain $\times_{k \in \mathcal{K}} [0, \phi^{max}_k]$. ∎

Theorem 1 shows that any equilibrium in terms of $\boldsymbol{\phi}$ is a fixed-point solution of Equation (10) in a vector form and the existence of such a solution is always guaranteed. Although the uniqueness of equilibrium cannot be guaranteed, it provides a sufficient condition on the function $\Phi(\cdot)$ of the route-level rates and the matrix $H(\cdot)$ of the effective gains in the compact domain of possible input rates to the resources, because each ϕ^{max}_k defines the maximum input rate to resource k when the delays and gains along any route equals 0 and 1, respectively.

In the next two sections, we will discuss applications of the equilibrium framework to understand the impacts of the strategic behaviors of various parties, e.g., peering and routing decisions of CPs and capacity planning decisions of ISPs.

4 LINEAR LOSSY NETWORKS

In this section, we consider systems that consist of independent routes among which no resource is shared, i.e., $l \cap l' = \emptyset$ for any $l, l' \in \mathcal{L}$. Despite being a special case, it represents important real scenarios of service differentiation and pricing, under which ISPs dedicate isolated network resources to form service classes and CPs choose routes in the form of peering decisions, e.g., premium peering [4, 18], based on pricing.

When routes are independet, we can focus on a typical route which forms a linear topology of K resources from CPs to end-users as shown in Figure 6. We drop the subscript l and denote the number of end-users of CP i by s_i. We also denote the aggregate sending rate and throughput of the route under study by φ and τ. $\Phi(\boldsymbol{\phi})$ becomes a scalar function defined as

$$\Phi(\boldsymbol{\phi}) = \sum_{i \in \mathcal{I}} s_i \Lambda_i \left(\sum_{\kappa=1}^{K} \Psi_\kappa(\phi_\kappa, \mu_\kappa), \prod_{\kappa=1}^{K} \Psi^g_\kappa(\phi_\kappa, \mu_\kappa) \right).$$

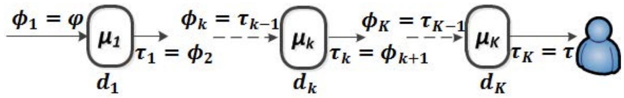

Figure 6: An illustration of an independent route with a linear topology.

We define the output rate of resource k as $\tau_k = \phi_k \Psi_k^g(\phi_k)$. Consequently, $\tau = \tau_K$ and the input rates to the individual resources satisfy $\phi_1 = \varphi$ and $\phi_{k+1} = \tau_k$ for all $k > 1$.

Corollary 1. *When $K = 1$, there is always a unique equilibrium. When $K = 2$, the equilibrium is unique if*

$$g_1 \frac{\partial \Phi}{\partial \phi_1} + \frac{\partial \Psi_1^g}{\partial \phi_1} \Phi(\boldsymbol{\phi}) \neq \left(\frac{\partial \Phi}{\partial \phi_1} - 1 \right) \left[g_1 - \left(\frac{\partial \Phi}{\partial \phi_2} \right)^{-1} \right] \tag{12}$$

for all $(\phi_1, \phi_2) \in [0, \phi_1^{max}] \times [0, \phi_2^{max}]$, where $g_1 = \Psi_1^g(\phi_1)$

$$\text{and} \quad \frac{\partial \Phi}{\partial \phi_k} = \sum_{i \in \mathcal{I}} s_i \left[\frac{\partial \Lambda_i}{\partial D} \frac{\partial \Psi_k}{\partial \phi_k} + \frac{\partial \Lambda_i}{\partial G} \frac{G}{g_k} \frac{\partial \Psi_k^g}{\partial \phi_k} \right]. \tag{13}$$

Proof of Corollary 1: For $K = 1$, $H(\boldsymbol{\phi}) = 1$ and $\Phi(\boldsymbol{\phi}) = \Phi(\phi_1) = \sum_{i \in \mathcal{I}} s_i \Lambda_i(\Psi_1(\phi_1, \mu_1), \Psi_1^g(\phi_1, \mu_1))$ and therefore,

$$\nabla_{\boldsymbol{\phi}} \Phi(\boldsymbol{\phi}) = \sum_{i \in \mathcal{I}} s_i \left(\frac{\partial \Lambda_i}{\partial D} \frac{\partial \Psi_1}{\partial \phi_1} + \frac{\partial \Lambda_i}{\partial G} \frac{\partial \Psi_1^g}{\partial \phi_1} \right).$$

By Assumption 1 and 2, $\partial \Lambda_i / \partial D < 0$, $\partial \Lambda_i / \partial G \geq 0$, $\partial \Psi_1 / \partial \phi_1 \geq 0$ and $\partial \Psi_1^g / \partial \phi_1 \leq 0$, implying $\nabla_{\boldsymbol{\phi}} \Phi(\boldsymbol{\phi}) \leq 0$. Consequently, condition (11) becomes $|\nabla_{\boldsymbol{\phi}} \Phi(\boldsymbol{\phi}) - 1| \neq 0$, which always holds. For $K = 2$, $H(\boldsymbol{\phi}) = (1, \Psi_1^g(\phi_1, \mu_1)) = (1, g_1)$ and the sufficient condition (11) becomes

$$\left\| \begin{bmatrix} \partial \Phi / \partial \phi_1 \\ \partial \Phi / \partial \phi_2 \end{bmatrix} [1, g_1] + \begin{bmatrix} 0 & \partial \Psi_1^g / \partial \phi_1 \\ 0 & 0 \end{bmatrix} - \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right\| \neq 0,$$

which can be further simplified to the condition (12). ∎

Corollary 1 states that if a route consists of a single resource, the uniqueness of equilibrium is guaranteed; however, when $K = 2$, the sufficient condition of (11) is reduced to (12). Notice that the left hand side of (12) can be written as $\partial \tau_1 / \partial \phi_1$ or $\partial \phi_2 / \partial \phi_1$, which captures the changes in the input rate ϕ_2 when ϕ_1 changes. Since $\partial \Phi(\boldsymbol{\phi}) / \partial \phi_k$ is negative by Assumption 1 and 2, the right hand side of (12) is always negative; and therefore, condition (12) is satisfied if $\partial \tau_1 / \partial \phi_1$ is non-negative. Although not assumed, it is a natural consequence if network resources are controlled in a *work conserving* manner.

Assumption 3 (Work Conserving). *Given any required delay d, resource k can accommodate a maximum throughput $T_k(d)$, and $T_k(d)$ is continuous and non-decreasing in d. Under any input rate λ_k, the induced throughput satisfies*

$$\tau_k \triangleq \lambda_k \Psi_k^g(\lambda_k, \mu_k) = \min \{\lambda_k, T_k(\Psi_k(\lambda_k, \mu_k))\}.$$

Lemma 1. *Under Assumption 3, Assumption 1 implies that throughput $\tau_k = \lambda_k \Psi_k^g(\lambda_k, \mu_k)$ is non-decreasing in λ_k.*

Proof of Lemma 1: By Assumption 1, i.e., $\Psi_k(\lambda_k)$ is non-decreasing in λ_k, and by Assumption 3, $T_k(d)$ is non-decreasing in d, we can deduce the composite function $T_k(\Psi_k(\lambda_k))$ is non-decreasing in λ_k. Because under Assumption 3, $\lambda_k \Psi_k^g(\lambda_k) = \min\{\lambda_k, T_k(\Psi_k(\lambda_k))\}$, the throughput $\lambda_k \Psi_k^g(\lambda_k)$ must be non-decreasing in λ_k as well. ∎

Assumption 3 states that if resource k can accommodate a throughput rate $T_k(\Psi_k(\lambda))$ under the induced delay $\Psi_k(\lambda)$, it will fully utilize the capacity to accommodate $T_k(\Psi_k(\lambda))$, unless the demand λ_k is less. Consequently, Lemma 1 shows that under this work conservation assumption, the monotonicity of throughput τ_k can be guaranteed as a result of Assumption 1.

Theorem 2. *Under Assumption 1 to 3, there always exists a unique equilibrium for any system with a linear topology.*

Proof of Theorem 2: Because any input rate ϕ_k determines the delay d_k, gain g_k and throughput τ_k, and under a linear topology $\phi_1 = \varphi$ and $\phi_k = \tau_{k-1}$ for $k > 1$, the sending rate φ uniquely determines all other parameters in the system. Suppose we have two different rates $\varphi^* > \varphi^\star$ under equilibrium, we prove the uniqueness of equilibrium by contradiction.

As $\varphi^* > \varphi^\star$, $\phi_1^* > \phi_1^\star$ and by Assumption 1, $d_1^* \geq d_1^\star$ and $g_1^* \leq g_1^\star$. Furthermore, with Assumption 3 and Lemma 1, we deduce that $\tau_1^* > \tau_1^\star$. Because $\phi_k = \tau_{k-1}$ for $k > 1$, the same logic can be applied to derive $d_k^* \geq d_k^\star$, $g_k^* \leq g_k^\star$ and $\tau_k^* > \tau_k^\star$ for any resource k. Consequently, we deduce that the aggregate delay and gain satisfy $D^* \geq D^\star$ and $G^* \leq G^\star$. This contradicts Assumption 2, which states that the demand will not increase under a higher delay and a lower gain. ∎

Theorem 2 shows that under an assumption of work conservation, any linear topology obtains a unique equilibrium. Thus, we denote the unique equilibrium by $(\boldsymbol{\phi}^*, \boldsymbol{d}^*, \boldsymbol{g}^*, \varphi^*, D^*, G^*)$. Next, we analyze the impact of ISPs' capacities and CPs' user population on the throughput τ^* of the route and the delays $\boldsymbol{d}^*$ and gains $\boldsymbol{g}^*$ of the individual resource under the equilibrium. In particular, we denote the capacities of resources by $\boldsymbol{\mu}$ and the numbers of end-users of the CPs by $\boldsymbol{s}$, and express the equilibrium as functions of $\boldsymbol{s}$ and $\boldsymbol{\mu}$ such as $\phi_k^*(\boldsymbol{\mu})$ and $\phi_k^*(\boldsymbol{s})$.

Theorem 3. *Let $\boldsymbol{s}' = (s_i', \boldsymbol{s}_{-i})$ and $\boldsymbol{s} = (s_i, \boldsymbol{s}_{-i})$ for any vector $\boldsymbol{s}_{-i}$ of end-users of CPs other than i with the condition $s_i' > s_i$. We must have $\varphi^*(\boldsymbol{s}') > \varphi^*(\boldsymbol{s})$ and for any resource k,*

$$\tau_k^*(\boldsymbol{s}') \geq \tau_k^*(\boldsymbol{s}), \ \ d_k^*(\boldsymbol{s}') \geq d_k^*(\boldsymbol{s}) \ \ \text{and} \ \ g_k^*(\boldsymbol{s}') \leq g_k^*(\boldsymbol{s}).$$

Let $\boldsymbol{\mu}' = (\mu_k', \boldsymbol{\mu}_{-k})$ and $\boldsymbol{\mu} = (\mu_k, \boldsymbol{\mu}_{-k})$ for any $\boldsymbol{\mu}_{-k}$ and $\mu_k' > \mu_k$. We must have $\tau^*(\boldsymbol{s}') \geq \tau^*(\boldsymbol{s})$ and for any $\kappa > k$,

$$\phi_\kappa^*(\boldsymbol{\mu}') \geq \phi_\kappa^*(\boldsymbol{\mu}), \ \ d_\kappa^*(\boldsymbol{\mu}') \geq d_\kappa^*(\boldsymbol{\mu}) \ \ \text{and} \ \ g_\kappa^*(\boldsymbol{\mu}') \leq g_\kappa^*(\boldsymbol{\mu}).$$

Proof of Theorem 3: We first show that $\varphi^*(\boldsymbol{s}') > \varphi^*(\boldsymbol{s})$ by contradiction. Suppose $\varphi^*(\boldsymbol{s}') \leq \varphi^*(\boldsymbol{s})$, because $\phi_1 = \varphi$ and by Assumption 3, we deduce $d_k^*(\boldsymbol{s}') \leq d_k^*(\boldsymbol{s})$ and $g_k^*(\boldsymbol{s}') \geq g_k^*(\boldsymbol{s})$ for any resource k, which implies $D^*(\boldsymbol{s}') \leq D^*(\boldsymbol{s})$ and $G^*(\boldsymbol{s}') \geq G^*(\boldsymbol{s})$. However, under no worse congestion and by Assumption 2, each rate Λ_i will be non-decreasing under $\boldsymbol{s}'$ and thus $\varphi^*(\boldsymbol{s}') > \varphi^*(\boldsymbol{s})$, which reaches a contradiction. Given $\varphi^*(\boldsymbol{s}') > \varphi^*(\boldsymbol{s})$, by Assumption 1 and Lemma 1, we can further deduce that $d_k^*(\boldsymbol{s}') \geq d_k^*(\boldsymbol{s})$, $g_k^*(\boldsymbol{s}') \leq g_k^*(\boldsymbol{s})$ and $\tau_k^*(\boldsymbol{s}') \geq \tau_k^*(\boldsymbol{s})$ for any resource k.

We then show that $\tau_k^*(\boldsymbol{\mu}') \geq \tau_k^*(\boldsymbol{\mu})$ by contradiction. Suppose $\tau_k^*(\boldsymbol{\mu}') < \tau_k^*(\boldsymbol{\mu})$, by Assumption 3, $d_\kappa^*(\boldsymbol{s}') \leq d_\kappa^*(\boldsymbol{s})$ and $g_\kappa^*(\boldsymbol{s}') \geq g_\kappa^*(\boldsymbol{s})$ for all downstream resources $\kappa > k$. This is also true for resource k and all upstream resources due to the monotonicity of throughput of Lemma 1 and the capacity impact in Assumption 1, which implies that $\phi_\kappa^*(\boldsymbol{\mu}') \leq \phi_\kappa^*(\boldsymbol{\mu})$ for all upstream resources $\kappa < k$. Similarly as before, this implies that $D^*(\boldsymbol{\mu}') \leq D^*(\boldsymbol{\mu})$ and $G^*(\boldsymbol{\mu}') \geq G^*(\boldsymbol{\mu})$, and therefore, $\varphi^*(\boldsymbol{\mu}') > \varphi^*(\boldsymbol{\mu})$. However, since

$\phi_1 = \varphi$, this contradicts with $\phi_1^*(\boldsymbol{\mu}') \leq \phi_1^*(\boldsymbol{\mu})$. Given $\tau_k^*(\boldsymbol{\mu}') \geq \tau_k^*(\boldsymbol{\mu})$, by Assumption 1 and Lemma 1, we further deduce $d_\kappa^*(\boldsymbol{\mu}') \geq d_\kappa^*(\boldsymbol{\mu})$, $g_\kappa^*(\boldsymbol{\mu}') \leq g_\kappa^*(\boldsymbol{\mu})$ and $\phi_\kappa^*(\boldsymbol{\mu}') \geq \phi_\kappa^*(\boldsymbol{\mu})$ for all downstream resources $\kappa > k$. ∎

Theorem 3 shows the unilateral impact of population s_i and capacity μ_k on the equilibrium. When s_i increases, the aggregate sending rate φ^* must increase and no resource will induce a higher gain or a lower throughput or delay. This implies the route delay D and gain G will be non-decreasing and non-increasing, respectively, which further implies that the aggregate sending rate of the CPs other than CP i will be non-increasing, but that of CP i will increase. When μ_k increases, the aggregate throughput τ^* will not decrease. In particular, no downstream resource $\kappa > k$ will receive a lower input rate and induce a lower delay or a higher gain. However, if the decrease of delay and increase of gain at resource k cannot compensate the reversed effects at the subsequent downstream resources, the sending rate φ^* might decrease and relieve the congestion at the upstream resources $\kappa < k$ in the new equilibrium.

Implications: The above result provides some implications on the peering relationship between CPs and ISPs. In the context of Netflix-Comcast dispute [7], Netflix accused Comcast for throttling end-users' throughput at the last-mile. Our result shows that even Comcast provides isolated resources for Netflix, the end-to-end performance will degrade as Netflix becomes popular and attracts more users. Furthermore, even under a fixed user population, as Netflix increases its capacities to provide better service quality, it will increase the end-to-end demand from users and make downstream resources at Comcast more congested. Although the aggregate throughput of Netflix increases, the sending rates of some users might decrease, because the end-to-end performance in terms of either delay D or drop rate $1 - G$ (but not both) might degrade.

5 GENERAL LOSSLESS NETWORKS

In this section, we consider general topologies, but focus on lossless networks, i.e., $g_k = 1$ and $\phi_k = \tau_k$ for all $k \in \mathcal{K}$. We refer to both the sending rate and throughput by $\boldsymbol{\phi}$ and $\boldsymbol{\varphi}$. These scenarios model the cases where application protocols and end-users are very sensitive to packet losses and will reduce demand under losses, e.g., adapting to low-resolution for video streaming, such that losses rarely occur under equilibria.

We define a $|\mathcal{K}| \times |\mathcal{L}|$ routing matrix R with each entry $R_{kl} = \mathbb{1}_{\{k\in l\}}$. Notice that $R = H(\mathbf{0})^T$ and determines the network topology, since $\Psi_k^g(0) = 1$ holds naturally for any resource. Consequently, the network equilibrium of Definition 1 can be compactly represented in a vector form as

$$\boldsymbol{d} = \Psi(\boldsymbol{\phi}, \boldsymbol{\mu}), \quad \boldsymbol{\varphi} = \Phi(\boldsymbol{D}, \boldsymbol{s}), \quad \boldsymbol{\phi} = R\boldsymbol{\varphi} \text{ and } \boldsymbol{D} = R^T\boldsymbol{d},$$

where $\Phi_l(D_l, \boldsymbol{s}_l) = \sum_{i\in\mathcal{I}} s_{il}\Lambda_i(D_l, 1)$ and $\boldsymbol{s}_l$ is the lth column of $\boldsymbol{s}$, which is the $|\mathcal{I}| \times |\mathcal{L}|$ matrix of user population.

Theorem 4. *Under Assumption 1 and 2, there always exists a unique equilibrium for any lossless system.*

Proof of Theorem 4: Theorem 1 shows the existence of equilibrium. Here, we show the uniqueness by contradiction. Suppose for any fixed $\boldsymbol{\mu}$ and $\boldsymbol{s}$, there exist two equilibria $(\boldsymbol{\phi}', \boldsymbol{d}', \boldsymbol{\varphi}', \boldsymbol{D}') \neq (\boldsymbol{\phi}, \boldsymbol{d}, \boldsymbol{\varphi}, \boldsymbol{D})$. Let us define the difference of the equilibria by

$$(\boldsymbol{\phi}_\delta, \boldsymbol{d}_\delta, \boldsymbol{\varphi}_\delta, \boldsymbol{D}_\delta) = (\boldsymbol{\phi}', \boldsymbol{d}', \boldsymbol{\varphi}', \boldsymbol{D}') - (\boldsymbol{\phi}, \boldsymbol{d}, \boldsymbol{\varphi}, \boldsymbol{D})$$

and without loss of generality, we order the resources and routes such that $\boldsymbol{d}_\delta = (\boldsymbol{d}_\delta^+, \boldsymbol{d}_\delta^-)$ and $\boldsymbol{\varphi}_\delta = (\boldsymbol{\varphi}_\delta^+, \boldsymbol{\varphi}_\delta^-)$, where $\boldsymbol{d}_\delta^+$ and $\boldsymbol{\varphi}_\delta^+$ consist of the positive differences and $\boldsymbol{d}_\delta^-$ and $\boldsymbol{\varphi}_\delta^-$ consist of the non-positive ones. By Assumption 1 and 2, we can write $\boldsymbol{\phi}_\delta = (\boldsymbol{\phi}_\delta^+, \boldsymbol{\phi}_\delta^-)$ and $\boldsymbol{D}_\delta = (\boldsymbol{D}_\delta^-, \boldsymbol{D}_\delta^+)$ correspondingly, and express

$$\boldsymbol{\phi} = R\boldsymbol{\varphi} \quad \text{as} \quad \begin{bmatrix} \boldsymbol{\phi}_\delta^+ \\ \boldsymbol{\phi}_\delta^- \end{bmatrix} = \begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} \boldsymbol{\varphi}_\delta^+ \\ \boldsymbol{\varphi}_\delta^- \end{bmatrix}.$$

By negating the sign of some components, we define

$$\tilde{R} = \begin{bmatrix} R_{11} & -R_{12} \\ R_{21} & -R_{22} \end{bmatrix},$$

$\tilde{\boldsymbol{\varphi}_\delta} = (\boldsymbol{\varphi}_\delta^+, -\boldsymbol{\varphi}_\delta^-)$ and $\tilde{\boldsymbol{D}_\delta} = (-\boldsymbol{D}_\delta^-, \boldsymbol{D}_\delta^+)$. By the linear equations $\boldsymbol{\phi} = R\boldsymbol{\varphi}$ and $\boldsymbol{D} = R^T\boldsymbol{d}$, we deduce

$$\tilde{R}\tilde{\boldsymbol{\varphi}_\delta} = \boldsymbol{\phi}_\delta \quad \text{and} \quad \tilde{\boldsymbol{\varphi}_\delta} \geq 0; \tag{14}$$

$$\tilde{R}^T(-\boldsymbol{d}_\delta) = \tilde{\boldsymbol{D}_\delta} \geq 0 \quad \text{and} \quad \boldsymbol{\phi}_\delta^T(-\boldsymbol{d}_\delta) < 0. \tag{15}$$

However, by using Farkas's Lemma [26], we conclude that the two conditions (14) and (15) cannot be satisfied simultaneously; and therefore, the equilibrium must be unique. ∎

Theorem 4 guarantees the uniqueness of network equilibrium for lossless systems. Similarly, we denote the unique equilibrium by $(\boldsymbol{\phi}^*, \boldsymbol{d}^*, \boldsymbol{\varphi}^*, \boldsymbol{D}^*)$ and express the equilibrium as a function of $\boldsymbol{\mu}$ and $\boldsymbol{s}$ such as $\boldsymbol{D}^*(\boldsymbol{s}, \boldsymbol{\mu})$ and $\boldsymbol{\varphi}^*(\boldsymbol{s}, \boldsymbol{\mu})$.

We study the impacts of capacities $\boldsymbol{\mu}$ and user population $\boldsymbol{s}$ on the route-level equilibrium $(\boldsymbol{\varphi}^*, \boldsymbol{D}^*)$ using primitives like $\nabla_{\boldsymbol{D}}\Phi$, $\nabla_{\boldsymbol{\mu}}\Psi$ and $\nabla_{\boldsymbol{\varphi}}D$. $\nabla_{\boldsymbol{D}}\Phi$ and $\nabla_{\boldsymbol{\mu}}\Psi$ are diagnal matrices with diagnal entries to be $\partial\Phi_l/\partial D_l$ and $\partial\Psi_k/\partial\mu_k$, respectively. Becuase $\boldsymbol{D} = R^T\boldsymbol{d} = R^T\Psi(\boldsymbol{\phi}, \boldsymbol{\mu}) = R^T\Psi(R\boldsymbol{\varphi}, \boldsymbol{\mu})$, by the chain rule, we can define $\nabla_{\boldsymbol{\varphi}}D \triangleq R^T\nabla_{\boldsymbol{\phi}}\Psi R$, where $\nabla_{\boldsymbol{\phi}}\Psi$ is a diagnal matrix with the kth diagnal entry being $\partial\Psi_k/\partial\phi_k$. Each entry $[\nabla_{\boldsymbol{\varphi}}D]_{l_1 l_2}$ captures the marginal delay on route l_2 due to the marginal change in the sending rate on route l_1, which equals the aggregate changes in the delays of resources $\sum_{k\in l_1\cap l_2}\partial\Psi_k/\partial\phi_k$ shared by routes l_1 and l_2.

Theorem 5. *The impact of capacities $\boldsymbol{\mu}$ on the route-level equilibrium $(\boldsymbol{\varphi}^*, \boldsymbol{D}^*)$ can be characterized by*

$$\begin{cases} \nabla_{\boldsymbol{\mu}}\boldsymbol{D}^* = \left(I - \nabla_{\boldsymbol{\varphi}}D\nabla_{\boldsymbol{D}}\Phi\right)^{-1} R^T\nabla_{\boldsymbol{\mu}}\Psi; \\ \nabla_{\boldsymbol{\mu}}\boldsymbol{\varphi}^* = \left(I - \nabla_{\boldsymbol{D}}\Phi\nabla_{\boldsymbol{\varphi}}D\right)^{-1} \nabla_{\boldsymbol{D}}\Phi R^T\nabla_{\boldsymbol{\mu}}\Psi. \end{cases} \tag{16}$$

The impact of users $\boldsymbol{s}$ on $(\boldsymbol{\varphi}^*, \boldsymbol{D}^*)$ can be characterized by

$$\begin{cases} \nabla_{\boldsymbol{s}}\boldsymbol{D}^* = \left(I - \nabla_{\boldsymbol{\varphi}}D\nabla_{\boldsymbol{D}}\Phi\right)^{-1} \nabla_{\boldsymbol{\varphi}}D\nabla_{\boldsymbol{s}}\Phi; \\ \nabla_{\boldsymbol{s}}\boldsymbol{\varphi}^* = \left(I - \nabla_{\boldsymbol{D}}\Phi\nabla_{\boldsymbol{\varphi}}D\right)^{-1} \nabla_{\boldsymbol{s}}\Phi. \end{cases} \tag{17}$$

Proof of Theorem 5: Given $\boldsymbol{D}^*(\boldsymbol{s}, \boldsymbol{\mu})$ is under equilibrium,

$$\boldsymbol{D}^*(\boldsymbol{s}, \boldsymbol{\mu}) = R^T\Psi(R\Phi(\boldsymbol{D}^*(\boldsymbol{s}, \boldsymbol{\mu}), \boldsymbol{s}), \boldsymbol{\mu}).$$

By differentiating $\boldsymbol{\mu}$ and $\boldsymbol{s}$ on both sides, we obtain

$$\begin{aligned} \nabla_{\boldsymbol{\mu}}\boldsymbol{D}^* &= R^T\left[\nabla_{\boldsymbol{\phi}}\Phi R\nabla_{\boldsymbol{D}}\Phi\nabla_{\boldsymbol{\mu}}\boldsymbol{D}^* + \nabla_{\boldsymbol{\mu}}\Phi\right] \\ &= R^T\nabla_{\boldsymbol{\phi}}\Phi R\nabla_{\boldsymbol{D}}\Phi\nabla_{\boldsymbol{\mu}}\boldsymbol{D}^* + R^T\nabla_{\boldsymbol{\mu}}\Phi \\ &= \nabla_{\boldsymbol{\varphi}}D\nabla_{\boldsymbol{D}}\Phi\nabla_{\boldsymbol{\mu}}\boldsymbol{D}^* + R^T\nabla_{\boldsymbol{\mu}}\Phi \end{aligned}$$

$$\begin{aligned}\nabla_s D^* &= R^T\nabla_{\phi}\Phi R[\nabla_D\Phi\nabla_s D^* + \nabla_s\Phi]\\ &= \nabla_{\varphi}D[\nabla_D\Phi\nabla_s D^* + \nabla_s\Phi]\\ &= \nabla_{\varphi}D\nabla_D\Phi\nabla_s D^* + \nabla_{\varphi}D\nabla_s\Phi\end{aligned}$$

By rearranging the above, we obtain

$$\begin{cases}(I - \nabla_{\varphi}D\nabla_D\Phi)\nabla_{\mu}D^* = R^T\nabla_{\mu}\Phi;\\ (I - \nabla_{\varphi}D\nabla_D\Phi)\nabla_s D^* = \nabla_{\varphi}D\nabla_s\Phi,\end{cases} \tag{18}$$

which implies the first equations of (16) and (17). Similarly, given $\varphi^*(s, \mu)$ is under equilibrium, we have

$$\varphi^*(s, \mu) = \Phi(R^T\Psi(R\varphi^*(s, \mu), \mu), s).$$

By differentiating μ and s on both sides and using the same logic, we deduce the second equations of (16) and (17). ■

Theorem 5 provides the sensitivity analysis and comparative statics for the route-level equilibrium (φ^*, D^*) under varying μ and s, where marginal changes in equilibrium are presented as functions of primitive metrics like the gradient of Ψ and Φ. Although more abundant capacities reduce delays ($\nabla_{\mu}\Psi \le \mathbf{0}$) and increase throughput ($\nabla_D\Phi \le \mathbf{0}$), Equation (16) shows that this impact on the equilibrium might not be monotonic on every route due to the inter-dependency among the routes, as the inverse of a positive matrix might not be positive definite. Similarly, although more users induce higher throughput ($\nabla_s\Phi \ge \mathbf{0}$) and delay ($\nabla_{\varphi}D \ge \mathbf{0}$), this monotonicity might not hold under the equilibrium as shown by Equation (17).

Although the network equilibrium is partially affected by s, this user population matrix is ultimately determined by the strategic routing decisions of CPs. Based on our equilibrium framework, we can further study the CPs' routing decisions. To model end-users in different geographical regions, we denote the set of geographical regions by $\mathcal{J}$ and the maximum number of users that are interested in using CP i in region j by m_{ij}. The last mile of each route targets certain region $j \in \mathcal{J}$; and therefore, we denote the set of routes to reach region j by $\mathcal{L}^j$ and define $\mathcal{L} = \cup_{j\in\mathcal{J}}\mathcal{L}^j$ as the set of all possible routes. We define the set of routes available for CP i to route traffic to the end-users in region j by $\mathcal{L}_i^j \subseteq \mathcal{L}^j$ and define $\mathcal{L}_i = \cup_{j\in\mathcal{J}}\mathcal{L}_i^j$ as the set of routes available to SP i. Consequently, we define $\mathcal{M} \triangleq \{m_{ij} : i \in \mathcal{I}, j \in \mathcal{J}\}$ as the aggregate user demand in the regions and $\mathcal{S}(\mathcal{M}) = \times_{i\in\mathcal{I}}\mathcal{S}_i(\mathcal{M})$ as the set of all feasible routing strategies of the CPs, where each $\mathcal{S}_i(\mathcal{M})$ is defined as

$$\mathcal{S}_i = \left\{s_i : s_{il} \ge 0, \textstyle\sum_{l\in\mathcal{L}_i^j} s_{il} = \sum_{l\in\mathcal{L}^j} s_{il} = m_{ij}, \forall j \in \mathcal{J}\right\},$$

which is a simplex of strategy space constrained by $s_{il} = 0$ for all $l \notin \mathcal{L}_i$, i.e., CPs cannot use any route they do not own.

Under any feasible routing profile $s \in \mathcal{S}$, the unique network equilibrium determines the end-to-end delay $D^*(s)$ along the routes. To optimize users' performance in terms of minimizing their delays, among all the available routes, CPs might always want to route users along the route with the minimum delay. However, CPs' routing strategies are inter-dependent, which leads to the definition of a Wardrop equilibrium [31].

Definition 2 (Wardrop Equilibrium). *A feasible routing strategy profile $s^* \in \mathcal{S}(\mathcal{M})$ is a Wardrop routing equilibrium if for any route $l \in \mathcal{L}$ and CP $i \in \mathcal{I}$ with $s^*_{il} > 0$,*

$$D_l^*(s^*) \le D_{l'}^*(s^*), \quad \forall l' \in \mathcal{L}_i^{J(l)},$$

where $J(l)$ denotes the region of destination of route l.

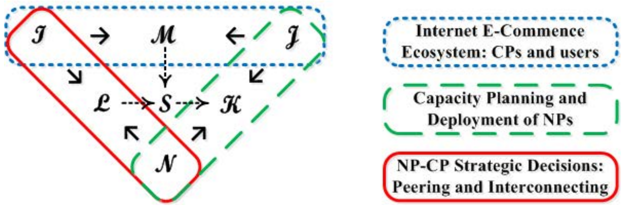

Figure 7: A macroscopic model of the Internet transport ecosystem.

Definition 2 states that under a Wardrop equilibrium, if any route $l \in \mathcal{L}^j$ is used by CP i to serve its users in region j, there does not exist any other feasible route $l' \in \mathcal{L}_i^j$ for CP i to serve users in region j whose end-to-end delay $D_{l'}^*$ is strictly smaller than the delay D_l^* of the original route l. Notice that this definition is built upon the delays $D^*(s)$ under the unique equilibrium as a function of the routing strategies of the CPs.

Theorem 6 ([29]). *A feasible routing strategy profile $s^* \in \mathcal{S}(\mathcal{M})$ is a Wardrop routing equilibrium if and only if*

$$(s - s^*)^T D^*(s^*) \ge \mathbf{0}, \quad \forall s \in \mathcal{S}(\mathcal{M}). \tag{19}$$

There always exists such a Wardrop routing equilibrium and the equilibrium is unique if D^* is strictly monotone [5], i.e.,

$$(D^*(s') - D^*(s))(s' - s) > \mathbf{0}, \ \forall s' \ne s, s', s \in \mathcal{S}(\mathcal{M}). \tag{20}$$

Theorem 6 characterizes the Wardrop equilibrium as a form of variational inequality [5] in (19). Since the feasible domain $\mathcal{S}(\mathcal{M})$ is compact, the existence of equilibrium can be guaranteed. Notice that a sufficient condition for the uniqueness of equilibrium, i.e., D^* being strict monotone, is that the gradient matrix $\nabla_s D^*$ defined in Equation (17) is positive definite, which implies that the increase in user demand along any route will not reduce the delay along any other route.

6 A BROADER INTERNET ECOSYSTEM MODEL AND APPLICATIONS

In general, we consider a set $\mathcal{N}$ of *network providers* (NPs), which include any type of ISP, e.g., transit and access, CDNs and any physical entity that owns network resources. The set of all network resources can be defined as $\mathcal{K} \triangleq \cup_{n\in\mathcal{N}}\mathcal{K}_n$, where each $\mathcal{K}_n$ defines the subset of resources owned by NP $n \in \mathcal{N}$. As a final result, a holistic Internet tranport ecosystem model can be described by a tuple $(\mathcal{I}, \mathcal{J}, \mathcal{K}, \mathcal{L}, \mathcal{M}, \mathcal{N})$.

Figure 7 illustrates the relationships among the different entities in the macroscopic Internet transport ecosystem model. The top blue rectangle shows the Internet E-commerce ecosystem where CPs $\mathcal{I}$ interact with users in various regions $\mathcal{J}$ and influence the demand $\mathcal{M}$ of users, which is not explicitly modeled in this work. This application layer of the Internet is enabled by the underlying physical transport system built by the NPs $\mathcal{N}$. The right green rectangle represents the supply-side where NPs make capacity planning decisions to deploy network resources $\mathcal{K}$ that form the routes $\mathcal{L}^j$ targeted for any region $j \in \mathcal{J}$. The left red rectangle represents the demand-side where NPs negotiate interconnection and peering contracts with CPs and other NPs. The interconnections form the physical Internet topology; while the peering agreements form the effective routes $\mathcal{L}_i$ used by each CP i's traffic. Under this macroscopic model, the user demand $\mathcal{M}$ and available routes $\mathcal{L}$ form the space $\mathcal{S}$ of routing strategies of the CPs, and any particular routing decision will eventually determine the delay and gain metrics on the resources $\mathcal{K}$ as well as the end-to-end delay D_l and gain G_l for each route $l \in \mathcal{L}$.

In the previous two sections, we demonstrated the use of our equilibrium framework for analyzing the impact of the capacity planning of ISPs and the routing decisions of CPs. Because the network equilibrium solution is built upon physical models of supply and demand of the Internet transport services, it can be used as a building block to understand the consequences of various strategic decisions of ASes. Consequently, higher-layer game-theoretical and optimization models can be established to study the business interactions among the ASes and desirable network protocols and policies for the Internet ecosystem. We briefly discuss some further applications of the macroscopic ecosystem model as follows.

Internet peering: The effects of peering determine the Internet topology and are reflected via an enriched set $\mathcal{L}$ of available routes. Furthermore, any NP n can control its resources $\mathcal{K}_n$ to create different routes, e.g., public and private peering points, and configure BGP export policies to make specific routes $\mathcal{L}_i$ available to any peering counter-party $i \in \mathcal{I}$ based on peering agreements such as premium peering. Thus, peering can be understood as a result of NPs' strategic controls of network resources and topology in the transport layer.

CDN and in-network caching: Although our framework does not explicitly model the storage capacities of the NPs, the use of CDN or caching can be reflected by the changes in the source ASes of contents. Because cached contents will have shorter routes towards end-users, CPs' decisions on cache deployment of CDNs [14] will effectively change the routing of contents towards users.

7 RELATED WORK

Early studies of network economics focused on the impact of selfish routing [21, 23, 27, 30] on network efficiency. Orda *et al.* [21] studied a routing game under which users split throughput demands among parallel links. Roughgarden and Tardos [27] analyzed the performance degeneration caused by selfish routing. These theoretical works assume fixed demands of users and only considered a single latency metric. We model elastic demands driven by network protocols such as TCP and characterizes both the drop rate and delay metrics. Teixeira et al. [30] showed via controlled experiments that intra-domain hot-potato routing causes high delays and slow convergence for inter-domian BGP [24] routes. Based on realistic topologies and traffic demands, Qiu *et al.* [23] studied selfish overlay routing in intra-domain environments via simulations. We focus on the macroscopic ecosystem where CPs make inter-domain routing decisions to enhance the QoS for end-users.

As the Internet topology is driven by the bilateral business relationship [6, 12] between ASes, many recent studies have focused on Internet peering [3, 4, 17–19]. Gao [12] characterized the *valley-free* property to infer ASes' business relationships based on the BGP routing protocol and data. Castro *et al.* [3] revealed the presence of remote peering, where remote networks peer via a layer-2 provider. Faratin *et al.* [6] and Lodhi *et al.* [17] discussed the complexity of peering and the emergence of new agreements, e.g., premium peering [4, 18], which however raised new peering disputes [7]. To resolve peering disputes, Ma *et al.* [19] designed a multilateral profit sharing mechanism for ISP settlements. Although CPs can obtain better QoS by using premium peering [4, 18] with access providers, the impact of such peering agreements on the resulting routing behaviors of CPs and the QoS of applications are largely unknown. Our equilibrium model captures the QoS in terms of both drop rates and delays of various routes, on top which Internet peering agreements can be better analyzed.

Extensive research was conducted to understand the QoS of TCP traffic flows under congestion control [11] and Active Queueing Management (AQM) [8] schemes. Mathis *et al.* [20] first proposed a renewal theory model for TCP Reno. Padhye *et al.* [22] derived the PFTK-formula that describes the TCP throughput as a function of loss rate and round trip time; however, both of which need to be known. Firoiu and Borden [8] analyzed the interactions between TCP and a bottleneck RED [11] queue using a fixed-point method. Bu and Towsley [2] extended the fixed-point framework for multiple bottlenecks, and Gibbens *et al.* [13] applied an M/M/1/B link model in a similar framework. The most extensive model was developed in Firoiu *et al.* [9] which consists of seven sets of equations and was evaluated by numerical methods; however, neither the existence nor the uniqueness of solution has been settled. Instead of modeling the detailed AQM and transport protocols, we apply general supply and demand functions to model network capacities and protocols. Consequently, we are able to derive the existence and uniqueness properties of equilibrium for general network topologies.

8 CONCLUSIONS

In this paper, we present a macroscopic network equilibrium model for the Internet transport ecosystem, which is built upon the supply of network capacity resources and the throughput demands driven by network protocols for transport services. Under such a network equilibrium, QoS metrics of drop rate and delay can be characterized for all the end-to-end routes. Through fixed-point analyses, we show the existence of equilibrium and its uniqueness under linear topologies or lossless scenarios. Through sensitivity analyses, we show the impacts of user demands and resource capacities on the network equilibrium, which provide implications of Netflix-Comcast type of peering on the QoS of users. Further studies of peering and caching can also be analyzed under our framework.

REFERENCES

[1] [n.d.]. The U.S. FCC's Open Internet Order. [Online] Available: https://www.fcc.gov/document/fcc-releases-open-internet-order.
[2] T. Bu and D. Towsley. 2001. Fixed Point Approximations for TCP behavior in an AQM Network. *ACM SIGMETRICS Performance Evaluation Review* 29, 1 (2001), 216–225.
[3] Ignacio Castro, Juan Camilo Cardona, Sergey Gorinsky, and Pierre Francois. 2014. Remote Peering: More Peering without Internet Flattening. In *Proceedings of the*

ACM CoNEXT Conference.

[4] Costas Courcoubetis, Laszlo Gyarmati, Nikolaos Laoutaris, Pablo Rodriguez, and Kostas Sdrolias. 2016. Negotiating Premium Peering Prices: A Quantitative Model with Applications. *ACM Transactions on Internet Technology (TOIT)* 16, 2 (April 2016).

[5] Francisco Facchinei and Jong-Shi Pang. 2003. *Finite-Dimensional Variational Ineqaulities and Complementarity Problems.* Springer-Verlag New York, Inc.

[6] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr. 2007. Complexity of Internet Interconnections: Technology, Incentives and Implications for Policy. *The 35th Research Conference on Communication, Information and Internet Policy (TPRC)* (2007).

[7] Nick Feamster. [n.d.]. Why Your Netflix Traffic is Slow, and Why the Open Internet Order Won't (Necessarily) Make It Faster. *[Online] Available: https://freedom-to-tinker.com/2015/03/25/why-your-netflix-traffic-is-slow-and-why-the-open-internet-order-wont-necessarily-make-it-faster/* ([n. d.]).

[8] V. Firoiu and M. Borden. 2000. A study of active queue management for congestion control. *In Proceedings of IEEE Infocom* (2000).

[9] Victor Firoiu, Jean-Yves Le Boudec, Don Towsley, and Zhi-Li Zhang. 2002. Theories and Models for Internet Quality of Service. *Proc. IEEE* 90, 9 (September 2002), 1565–1591.

[10] D. Fitzgerald. 2014. Netflix's share of Internet traffic grows. *The Wall Street Journal* (May 14 2014).

[11] S. Floyd and V. Jacobson. 1993. Random Early Detection Gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking* 1, 4 (August 1993).

[12] Linxin Gao. 2001. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking* 9, 6 (December 2001), 733–745.

[13] R. J. Gibbens, S.K. Sargood, C. Van Eijl, F. P. Kelly, H. Azmoodeh, R. N. Macfadyen, and N. W. Macfadyen. 2000. Fixed-Point Models for the End-to-End Performance Analysis of IP Networks. In *the 13th ITC Specialist Seminar: IP Traffic Measurement, Modeling and Management.* Monterey, California.

[14] Syed Hasan, Sergey Gorinsky, Constantine Dovrolis, and Ramesh K. Sitaraman. 2014. Trade-offs in Optimizing the Cache Deployments of CDNs. In *Proceedings of IEEE INFOCOM.* Toronto, ON, Canada.

[15] R. B. Kellogg. 1976. Uniqueness in the Schauder Fixed Point Theorem. *Proc. Amer. Math. Soc.* 60 (October 1976).

[16] R. Bruce Kellogg, Tien-Yien Li, and James A. Yorke. 1976. A constructive proof of the Brouwer fixed point theorem and computational results. *SIAM J. Numer. Anal.* 13, 4 (1976), 473–483.

[17] Aemen Lodhi, Nikolaos Laoutaris, Amogh Dhamdhere, and Constantine Dovrolis. 2015. Complexities in Internet peering: Understanding the "Black" in the "Black Art". In *Proceedings of IEEE INFOCOM.* Hong Kong.

[18] Richard T. B. Ma. 2017. Pay or Perish: The Economics of Premium Peering. *IEEE Journal on Selected Areas of Communications* 35, 2 (February 2017).

[19] Richard T. B. Ma, Dahming Chiu, John C.S. Lui, Vishal Misra, and Dan Rubenstein. 2011. On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers. *IEEE/ACM Transactions on Networking* 19, 3 (2011).

[20] M. Mathis, J. Semke, J. Mahdavi, and T. Ott. 1997. The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm. *ACM SIGCOMM Computer Communication Review* 27, 3 (1997), 67–82.

[21] Ariel Orda, Raphael Rom, and Nahum Shimkin. 1993. Competitive routing in multi-user communication networks. *IEEE/ACM Transactions on Networking* 1, 5 (October 1993).

[22] J. Padhye, V. Firoiu, D. F. Towsley, and J. F. Kurose. 2000. Modeling TCP Reno Performance: A Simple Model and Its Empirical Validation. *IEEE/ACM Transactions on Networking* 8, 2 (April 2000).

[23] Lili Qiu, Yang Richard Yang, Yin Zhang, and Scott Shenker. 2006. On selfish routing in Internet-like environments. *IEEE/ACM Transactions on Networking* 14, 4 (August 2006), 725–738.

[24] Y. Rekhter, T. Li, and S. Hares. 2006. A Border Gateway Protocol 4 (BGP-4). *RFC 4271* (2006).

[25] J.-C. Rochet and J. Tirole. 2003. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association* 1, 4 (2003), 990–1029.

[26] R. T. Rockafellar. 1979. *Convex Analysis.* Princeton University Press.

[27] Tim Roughgarden and Eva Tardos. 2000. How Bad is Selfish Routing?. In *IEEE Symposium on Foundations of Computer Science.* 93–102.

[28] Scott Shenker. 1995. Fundamental Design Issues for the Future Internet. *IEEE Journal on Selected Areas in Communication* 13, 7 (September 1995).

[29] M. J. Smith. 1979. The existence, uniqueness and stability of traffic equilibria. *Transportation Research Part B: Methodological* 13, 4 (December 1979), 295–304.

[30] Renata Teixeira, Aman Shaikh, Tim Griffin, and Jennifer Rexford. 2004. Dynamics of Hot-Potato Routing in IP Networks. *Proceedings of the ACM SIGMETRICS/Performance* (2004).

[31] John Glen Wardrop. 1952. Some Theoretical Aspects of Road Traffic Research. *ICE Proceedings: Engineering Divisions* 1, 5 (1952).

[32] Tim Wu. 2005. Network Neutrality, Broadband Discrimination. *Journal of Telecommunications and High Technology Law* 141 (2005).