

Pricing under Constraints in Access Networks: Revenue Maximization and Congestion Management

Prashanth Hande^{1,2}, Mung Chiang¹, Robert Calderbank¹, Junshan Zhang³

¹Department of Electrical Engineering, Princeton University, NJ 08544, USA

²Qualcomm Inc, NJ 08807, USA

³Arizona State University, AZ 85287, USA

Abstract—This paper investigates pricing of Internet connectivity services in the context of a monopoly ISP selling broadband access to consumers. We first study the optimal combination of flat-rate and usage-based access price components for maximization of ISP revenue, subject to a capacity constraint on the data-rate demand. Next, we consider time-varying consumer utilities for broadband data rates that can result in uneven demand for data-rate over time. Practical considerations limit the viability of altering prices over time to smoothen out the demanded data-rate. Despite such constraints on pricing, our analysis reveals that the ISP can retain the revenue by setting a low usage fee and dropping packets of consumer demanded data that exceed capacity. Regulatory attention on ISP congestion management discourages such “technical” practices and promotes economics based approaches. We characterize the loss in ISP revenue from an economics based approach. Regulatory requirements further impose limitations on price discrimination across consumers, and we derive the revenue loss to the ISP from such restrictions. We then develop partial recovery of revenue loss through non-linear pricing that does not explicitly discriminate across consumers. While determination of the access price is ultimately based on additional considerations beyond the scope of this paper, the analysis here can serve as a benchmark to structure access price in broadband access networks.

I. OVERVIEW

This paper studies the impact of access prices on the congestion management practices and revenue of a monopoly ISP, operating a single bottleneck link with fixed capacity. Although Internet data flows along multiple links on a route between source and destination, the end-user access link is typically the most constrained for capacity, and the major contributor to the connectivity price. Consumer data rate allocation can be determined by socially optimal prices in a competitive market on the one hand, or the revenue maximizing prices in a monopoly ISP market on the other hand. The analysis in this paper explores the latter and represents a benchmark: the most favorable outcome to the ISP and possibly the least favorable outcome to the consumers.

The analysis will serve as a basis to understand the pricing of shared access links in wireline and wireless broadband. The analysis has particular significance to pricing wireless broadband as consumers demand high-rate wireline-like applications and content over relatively lower-rate wireless broadband.

Access pricing is typically in the form of a flat rate that is independent of usage, or a usage based price, or some combination of the two pricing schemes [1], [2]. We quantify that a significant component of the monopoly ISP revenue is from *flat price if consumer price sensitivity is low* and

through *usage price if consumer price sensitivity is high*. Flat pricing is generally considered as the preferred choice of consumers [3], but our analysis indicates that flat pricing can lead to a significant loss of consumer net-utility, particularly when the consumers have low price sensitivity.

Consumer demand for data changes over hours of the day and days of the week, resulting in peak usage of networks that can be significantly high compared to average usage. Access ISPs face a mismatch between their revenue from average usage and cost incurred from peak usage of networks. Considerations on billing management and price simplicity discourage frequent changes in prices over time [3]. This limitation on ISP’s ability to manage peak aggregate demand through price variations can result in potential loss of revenue. Our analysis reveals that, despite the lack of flexibility to alter the time-dependent consumption of consumers through price variations, the ISP can *retain* the revenue through congestion management by *dropping packets* of consumer demanded data that *exceed* available capacity. We quantify the intuition that the revenue retention can be achieved through a combination of *low usage fee* that ensures sufficient consumer demand at all times, and a *high flat fee* that captures the remaining consumer net utility from the served data rates.

However, ISPs face regulatory hurdles [4], [5], including “network neutrality” concerns [6], that do not encourage congestion management by selectively dropping packets of consumer demanded data. The recent FCC notice [4] proposes draft language to codify a principle that would require a broadband ISP to treat lawful content, applications, and services in a nondiscriminatory manner. In a related decision, the CRTC recognizes [5] that “economic practices are the most transparent Internet traffic management practices”. It further notes that such economics based congestion management practices “match consumer usage with willingness to pay, thus putting users in control and allowing market forces to work”.

This suggests an alternative approach for the ISP: set sufficiently high access prices that ensure that the peak aggregate demand does not exceed the available capacity. In this case, the lack of flexibility to alter consumer data rate consumption through price variations results in a revenue loss. We show that the revenue loss is due to the inability of the ISP to *adapt the usage price to maximize capacity usage* and quantify the dependency of the revenue loss on consumer price sensitivity by showing that the *loss is higher if the consumers are highly price sensitive*. We study the pricing of a less than best-effort

connectivity service that the ISP can offer over the unused capacity to mitigate revenue loss.

Regulatory attention further discourages ISPs from practicing price discrimination across consumers. We show that the revenue loss from imposing uniform pricing across flows is due to the *inability of the ISP to charge the consumer-dependent flat price to capture the consumer net-utility*. Further, we quantify the dependency of the revenue loss on consumer price sensitivity and show that the *loss is higher if consumers are less price sensitive*. The revenue loss can be mitigated if the ISP can offer prices non-linear in data-volume or data-rate consumption, with discount on higher data-rate demand. The choice of the data consumption is left to the consumer, thus preventing explicit price discrimination among consumers. This problem is studied as ‘second-degree price discrimination’ in economics. We use related methods to quantify the loss recovery through non-linear pricing and show that the recovery is *less when consumer price sensitivity is low*.

The pricing analyses in the paper are applicable to real scenarios in broadband pricing, and we list below two examples:

Example 1: An ISP currently offers fixed monthly prices for broadband. With increasing broadband demand, the ISP notices traffic peaks during evenings in residential areas, and during mid-day in commercial locations. One option for the ISP to address the likely deterioration in customer experience is to spread the demand more evenly by charging less during non-peak times. The ISP can work out the additional cost of billing and logistics but is unsure how to price the non-peak traffic and if the revenue gain can offset the added billing costs. A second option is for the ISP to support, during non-peak hours, bulk traffic that can tolerate large delays. The revenue gain from this bulk-traffic may be used for capacity expansion to alleviate the customer experience problems during peak times.

Example 2: An ISP offers broadband services at a uniform price to all customers but notices significant variations in customer valuation of their services. One customer, a small-business owner, relies on her broadband to conduct business, and is willing to pay more to retain the ISP services. A different customer is an occasional user of broadband, and does not consider the service worth the current price. One option for the ISP is to explicitly price the two customers differently, but realizes that such discriminatory pricing could draw regulatory attention. Another option for the ISP is to implicitly price the two customers differently by offering non-linear pricing, with the customers choosing the consumed quantity.

A number of access network pricing models have been proposed, and an overview of the various models can be found in [7]. In economics, the two-part pricing model in [8] and the literature on club goods [9] have studied flat fee (membership fee) and usage based fee extensively. The work in [1] draws upon this literature and studies congestion pricing in networks. It is shown that an ISP with market power (*e.g.* monopoly ISP) can extract all of the net-utility from the consumers through a combination of flat and usage based pricing, resulting in data-rate allocation and overall welfare similar to the case of a competitive ISP market (the social optimum allocation). We derive an equivalent result for our model that serves as the

basis for further analysis of restrictions on the access prices. Revenue maximization is considered under a combination of access price and congestion price with fixed utility functions in [10] and the ISP incentive to increase capacity is studied. We analyze congestion management through access price under general utility functions for fixed capacity. Time-dependent pricing is considered in [11], where the authors propose a model with time-dependent consumer utilities and analyze the revenue loss to the monopoly ISP due to insufficient information about the consumer utilities. We assume complete information about consumer utilities and analyze the revenue loss to the ISP due to constraints on time variations of prices.

The rest of the paper is organized as follows. Section II presents the setup for the revenue maximization problems, based on which we study in Section III the optimal choice of flat and usage based components in access pricing. Section IV imposes constraints on price variations in time and quantifies the resulting revenue loss. Section V derives the revenue loss from imposition of the requirement of uniform pricing across flows and develops recovery of loss through non-linear pricing. We conclude with section VI.

II. SYSTEM SETUP AND BASIC NOTATION

The congestion constraint faced by an ISP is on the peak aggregate consumer data-rate. In contrast, the retail access price is based on the volume of data (measured in megabytes or MB), consumed over a specified time period T (typically a month). The pricing on data volume is equivalent to pricing on average data rate over time T . Therefore, ISPs face a mismatch, where revenue is accrued on average data rate but congestion cost is incurred on the peak data rate. To address this disparity, we note that the difference between peak data rate and the average data rate is reduced when measured over smaller time periods. Consider a unit time interval that is sufficiently small so that the peak data rate demand of a consumer in that time interval is a close approximation to the average data rate for that consumer in that interval¹.

Let $f \in \mathcal{F}$ be a consumer data flow with \mathcal{F} denoting the set of all flows. Let data rate for flow f in the interval $[(t-1), t]$ be given by x_f^t . The data volume consumption over time T is then given by $x_f = \sum_{t=1}^T x_f^t$, and the capacity constraint applies at every time instant t :

$$\sum_f x_f^t \leq C, \quad \forall t$$

A qualitative observation is that the shape of the utility function depends on the response of the content or application to varying data-rates, and the utility level represents the consumer’s need for the application or content. This motivates us to assume that the consumer’s utility level varies in time, but the shape of the utility function does not. The observation will have to be verified through analysis of real data. Let $\sigma_f^t u_f(x_f^t)$ be the utility to a consumer associated with flow f at time instant

¹In practice, measuring data usage in very short time intervals might not be amenable due to high implementation costs. It suffices to use an interval over which the peak usage can be estimated as a linearly scaled value of average usage with high confidence.

t , with factor σ_f^t denoting the time dependency of consumer's utility level.

We assume that the revenue maximizing ISP, through network measurements, has complete knowledge of the either the consumer utility function parameters or the probabilistic distribution of the parameters. A typical ISP practice is to collect time snap-shots of aggregate data-rate demand, which motivates us to assume a deterministic model of time variations in the utility level in section IV. Measuring flow-dependent parameters is more difficult, and we assume a probabilistic model of the utility level variation across consumers in section V. We note that the analysis in section IV can be easily extended to a probabilistic model of the time variations of utility levels.

Faced with time-varying consumer utilities, the ISP can charge a time-dependent price for connectivity $r_f^t(x_f^t)$ as a function of the allocated data rate x_f^t . Consumers maximize the net utility for each flow f :

$$\begin{aligned} & \text{maximize} && \sigma_f^t u_f(x_f^t) - r_f^t(x_f^t) \\ & \text{variable} && x_f^t \end{aligned} \quad (1)$$

The consumer demand function is the data rate that maximizes the net-utility in (1). One form of the price with linear increase in data-rate is given by: $r_f^t(x_f^t) = g_f^t + h_f^t x_f^t$. The flat price g_f^t is fixed for the duration of the time interval, irrespective of the allocated data rate. The usage based component involves a price h_f^t per unit data consumption. The demand function for this form of the price can be shown to be given by:

$$y_f^t(g_f^t, h_f^t) = \begin{cases} u_f'^{-1}(h_f^t/\sigma_f^t) & \text{if } g_f^t \leq \sigma_f^t u_f(y_f^t) - h_f^t y_f^t, \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The condition $\sigma_f^t u_f(x_f^t) - g_f^t - h_f^t x_f^t \geq 0$ ensures that consumers have non-negative utility. To simplify notation, we often use $y_f^t(h_f^t) = u_f'^{-1}(h_f^t/\sigma_f^t)$ with the implicit assumption that the flat price is low enough to ensure non-negative consumer net-utility. We do not consider an explicit penalty to consumer utility, as considered in [10], [11], from congestion due to aggregate data-rate demand. However, we note that the such a penalty can be easily incorporated into the consumer utility function, in which case $y_f^t(g_f^t, h_f^t)$ is the best-response update at a Nash-equilibrium of consumer data-rate demands.

The elasticity of demand η_f^t is a standard measure of the sensitivity of the consumer demand to price fluctuations [20], and is defined as

$$\eta_f^t = -\frac{\partial y_f^t}{\partial h_f^t} \frac{h_f^t}{y_f^t} \quad (3)$$

Often, we specialize the utility function to the standard alpha-fair utility form:

$$u_f(x_f) = \begin{cases} \log(x_f) & \text{if } \alpha_f = 1, \\ (1 - \alpha_f)^{-1} x_f^{1-\alpha_f} & \text{if } 0 \leq \alpha_f < 1 \end{cases} \quad (4)$$

for which we have $\eta_f^t = 1/\alpha_f$: the price sensitivity is inversely proportional to the parameter α_f and independent of utility level σ_f^t .

III. BASIC PRICE STRUCTURES

We first analyze the structure of the price $r_f^t(x_f^t) = g_f^t + h_f^t x_f^t$ that allows the ISP to manage congestion while maximizing revenue.

A. Price Structure

The revenue maximization problem for the monopoly ISP can be defined by the following:

$$\begin{aligned} & \text{maximize} && \sum_t \sum_f (g_f^t + h_f^t x_f^t) \\ & \text{subject to} && \sum_f x_f^t \leq C, \forall t \\ & && x_f^t \leq u_f'^{-1}(h_f^t/\sigma_f^t) \\ & && \sigma_f^t u_f(x_f^t) - g_f^t - h_f^t x_f^t \geq 0, \forall t \\ & \text{variables} && \{g_f^t, h_f^t, x_f^t\} \end{aligned} \quad (5)$$

Obviously, the revenue increases with higher flat fee component g_f^t , which can be set so that the consumer net utility is zero. The revenue from each flow is then $g_f^t + h_f^t x_f^t = \sigma_f^t u_f(x_f^t)$, which can be realized by any combination of flat and usage fee that can support a data-rate of x_f^t . If the usage fee h_f^t is such that the consumer demand $y_f^t(h_f^t)$ is strictly greater than the ISP provisioned data rate x_f^t , then flow data packets have to be dropped. However, the ISP can avoid packet drops by setting a sufficiently high usage price to reduce the consumer demand so that the aggregate demand is within the available capacity. It follows that $x_f^t = y_f^t(h_f^t)$ and the ISP revenue maximization problem (5) can be re-written as:

$$\begin{aligned} & \text{maximize} && \sum_t \sum_f \sigma_f^t u_f(y_f^t(h_f^t)) \\ & \text{subject to} && \sum_f y_f^t(h_f^t) \leq C, \forall t \\ & \text{variables} && \{h_f^t\} \end{aligned} \quad (6)$$

Let the maximized revenue in problems (5) and (6) from unrestricted pricing be R_u^* , which will be contrasted with maximized revenue under restrictions on pricing in later sections. Problem (6) can be easily decomposed into sub-problems at each time instant t , and has a solution given by the following

Theorem 1: An optimal pricing scheme that achieves the maximum in (5) is given by setting for each t :

$$\begin{aligned} h^{t*} &= \mu^t \\ x_f^{t*} &= u_f'^{-1}(\mu^t/\sigma_f^t) \\ \sum_f x_f^{t*} &= C \\ g_f^{t*} &= \sigma_f^t u_f(x_f^{t*}) - \mu^t x_f^{t*} \end{aligned} \quad (7)$$

The proof follows directly from the observation that the optimal price structure in (7) represents the KKT conditions [21] for the decomposed sub-problems of the optimization problem in (6). The optimal usage fee h^{t*} is the time dependent congestion price μ^t , which is the same across all flows f , and the optimal flat fee $g_f^{t*} = \sigma_f^t u_f(x_f^{t*}) - \mu^t x_f^{t*}$ is flow dependent, allowing the ISP to fully extract the consumer net-utility.

Let R_f^* be the revenue from flat component of the price and R_s^* the revenue from usage component so that $R_u^* = R_f^* + R_s^*$. It can be shown that

$$\frac{R_f^*}{R_s^*} = \frac{\sum_{f,t} \sigma_f^t u_f(x_f^{t*})}{\sum_{f,t} \sigma_f^t u_f'(x_f^{t*}) x_f^{t*}} - 1 \quad (8)$$

In an exemplary special case, we have the following result, whose proof is straightforward and is omitted.

Theorem 2: If utility functions are alpha-fair (4) with $\alpha_f = \alpha$ for all f , the ratio of flat revenue to usage dependent revenue is given by

$$\frac{R_f^*}{R_s^*} = \frac{\alpha}{1 - \alpha} \quad (9)$$

Remarks: Theorem 2 reveals that usage dependent revenue dominates with linear utilities ($\alpha \rightarrow 0$), while revenue from flat rate component dominates with log utilities ($\alpha \rightarrow 1$). It is apparent that the flat price is a significant component in the extraction of consumer net-utility, if the consumer price sensitivity is low.

B. An Illustrative Example

Consider a monopolist ISP providing connectivity service to 10 flows over an access link of capacity $C = 10$ Mbps. Consumers associated with each flow have utilities of the form given in (4) with $\alpha_f = \alpha$, and $\eta = 1/\alpha$ the common elasticity of demand across all flows. We generate the utility levels $\{\sigma_f^t\}$ as random variables uniformly distributed in $[\sigma_0, \sigma_1]$.

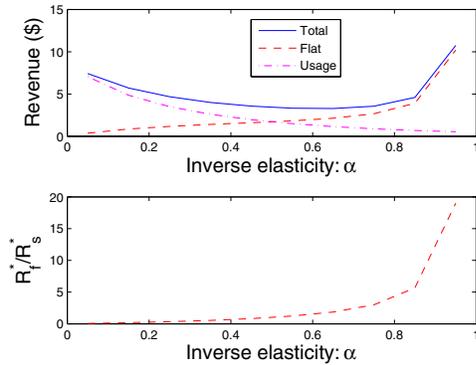


Fig. 1. Comparison of revenue from flat pricing to usage based pricing at different consumer demand elasticity.

Figure 1 illustrates the average revenue per unit time per unit flow generated by the monopolist ISP with no constraints on the prices. The flat component of the revenue, which enables the monopolist to completely extract the consumer net-utility, increases with decreasing elasticity of demand. The usage component of the revenue decreases with decreasing elasticity of demand. The lower part of Figure 1 plots the ratio of the flat component to the usage component of the revenue as given by (9), demonstrating the increased reliance on revenue from flat price at low consumer demand elasticity.

The ISP pricing flexibility, in practice, is restricted along time and across flows. In the following sections, we extend our model to quantify the revenue loss to the ISP and develop pricing schemes to mitigate the loss.

IV. TIME CONSTRAINED PRICING

As discussed in Section I, the common practice today is for the ISP to price the total data volume, or equivalently, the average data rate over the longer time horizon T . We model this by restricting time variations in $r_f^t(x_f^t) = g_f^t + h_f^t x_f^t$ and allow a single price $r_f = g_f + h_f x_f$ per flow across all times $t \in [1, \dots, T]$, with $x_f = \sum_t x_f^t$. The time-constrained ISP revenue maximization problem, whose optimum we denote by

R_t^* , can be written as:

$$\begin{aligned} & \text{maximize} && \sum_f (g_f + h_f \sum_t x_f^t) \\ & \text{subject to} && \sum_f x_f^t \leq C, \forall t \\ & && x_f^t \leq u_f'^{-1}(h_f/\sigma_f^t) \\ & && \sum_t \sigma_f^t u_f(x_f^t) - g_f - h_f \sum_t x_f^t \geq 0 \\ & \text{variables} && \{g_f, h_f, x_f^t\} \end{aligned} \quad (10)$$

The lack of flexibility in varying the price over time can potentially reduce the ISP revenue. One approach for the ISP is to set the time-independent usage fee h_f low enough to ensure that the aggregate consumer demand $\sum_f y_f^t(h_f)$ is higher than the available capacity C . The ISP can then choose to serve data rates $x_f^t \leq y_f^t(h_f)$ that maximizes the revenue. Consider the optimal usage prices in (7) for each time instant that was obtained as a solution to the revenue maximization problem in (5) without time constraints. The solution to the time-constrained pricing problem (10) can be obtained by first setting the usage price to the minimum across time of the optimal usage prices as given in (7): $h_f = \min\{h^{t*}\}$. Let the data rates served be given by $x_f^t = x_f^{t*}$, the optimal data rates served in (7) so that the capacity constraint is satisfied. If the time-independent flat fee g_f can be set such that

$$g_f = \sum_t \sigma_f^t u_f(x_f^{t*}) - \min\{h^{t*}\} \sum_t x_f^{t*},$$

the ISP can generate a revenue of $\sum_t \sigma_f^t u_f(x_f^{t*})$, which is the same as in the case without time constraints. This implies that $R_t^* = R_u^*$, with the time-constraint on pricing not impacting the ISP's revenue.

Note that the ISP could retain the revenue even with time-constraint on pricing by ensuring that consumer demands exceed capacity at all times. The downside is that ISP has to drop consumer data packets to limit the data rates to within the capacity. The control exerted by the ISP in packet-drop decisions with this approach has attracted regulatory attention [4] and is increasingly infeasible.

A. Time Constrained Pricing without Packet Dropping

Alternatively, the ISP can adopt economics based congestion management practices that are favored by some regulators [5]. One such approach is to set the usage price to $h_f = \max\{h^{t*}\}$, so that the worst case demand is within the capacity constraint. This ensures that the consumer aggregate demand is always within the capacity limit, thus avoiding the need for an explicit packet-drop decision by the ISP. However, capacity is not fully utilized at non-peak times, incurring a revenue loss. To characterize the revenue loss, consider the revenue maximization problem under this approach, with the corresponding optimal revenue denoted by R_p^* :

$$\begin{aligned} & \text{maximize} && \sum_f (g_f + h_f \sum_t y_f^t(h_f)) \\ & \text{subject to} && \sum_f y_f^t(h_f) \leq C, \forall t \\ & && \sum_t \sigma_f^t u_f(x_f^t) - g_f - h_f \sum_t x_f^t \geq 0 \\ & \text{variables} && \{g_f, h_f\} \end{aligned} \quad (11)$$

Intuitively, if utility levels $\{\sigma_f^t\}$ are well spread, then the revenue generated by the ISP will be close to that of an ISP with no restriction as in (5). A more likely pattern of data

usage is one where multiple flows overlap in time. The pricing then will be such that the data usage is within capacity limits during those times, and the capacity is wasted during other times, resulting in a loss of revenue compared to the case with no restriction.

To evaluate the effect of the spread in utility levels $\{\sigma_f^t\}$, consider flow-independent $\{\sigma^t\}$. Let $\sigma^m = \max\{\sigma^t\}$ be the maximum utility level for any flow over time. In an important special case, the ratio of the optimal revenue R_p^* to the revenue R_u^* without time-constrained pricing is given by the following theorem (proved in Appendix A).

Theorem 3: Assume utility functions are alpha-fair (4) with $\alpha_f = \alpha$ for all f and flow-independent utility levels $\{\sigma^t\}$. The ratio of optimal revenue R_p^* from problem (11) to the revenue R_u^* from problem (5) is given by

$$\frac{R_p^*}{R_u^*} = \frac{\sum_t \left(\frac{\sigma^t}{\sigma^m}\right)^{1/\alpha}}{\sum_t \left(\frac{\sigma^t}{\sigma^m}\right)} \quad (12)$$

Remarks: It is easy to see that $R_p^*/R_u^* \leq 1$. The revenue inefficiency stems from the inability of the ISP to charge the time-dependent usage fee to ensure full utilization of the capacity. The inefficiency is higher with lower values of α , which is equivalent to high consumer price sensitivity. This is consistent with the observation that a larger fraction of the revenue is dependent on the usage fee with lower α , as shown in Theorem 2.

We can develop a lower bound in terms of the minimum utility level $\sigma^n = \min\{\sigma^t\}$ by noting that

$$\begin{aligned} \frac{R_p^*}{R_u^*} &= \frac{\sum_t \left(\frac{\sigma^t}{\sigma^m}\right)^{1/\alpha}}{\sum_t \left(\frac{\sigma^t}{\sigma^m}\right)} \geq \frac{\sum_t \left(\frac{\sigma^n}{\sigma^m}\right)^{1/\alpha}}{\sum_t \left(\frac{\sigma^t}{\sigma^m}\right)} \\ &\geq \frac{T \left(\frac{\sigma^n}{\sigma^m}\right)^{1/\alpha}}{T \left(\frac{\sigma^m}{\sigma^m}\right)} \geq \left(\frac{\sigma^n}{\sigma^m}\right)^{1/\alpha} \end{aligned} \quad (13)$$

If σ^t can be modeled as realizations of a random variable, then we can invoke Jensen's inequality to specify a tighter bound.

$$\begin{aligned} \frac{R_p^*}{R_u^*} &= (\sigma^m)^{(1-1/\alpha)} \frac{E(\sigma^{1/\alpha})}{E(\sigma)} \\ &\geq (\sigma^m)^{(1-1/\alpha)} \frac{E(\sigma)^{1/\alpha}}{E(\sigma)} \\ &\geq \left(\frac{E(\sigma)}{\sigma^m}\right)^{(1/\alpha-1)} \end{aligned} \quad (14)$$

The bounds indicate that the revenue inefficiency is limited if the spread of utility levels is small. However, revenue can be highly inefficient if utility levels have a large spread, which is equivalent to consumers having similar and strong time preference for data consumption.

B. Pricing Scavenger Class for Revenue Loss Mitigation

When pricing is time-constrained and the packet-dropping approach is infeasible, it was demonstrated in Section IV-A that the capacity is underutilized during uncongested time. The underutilized capacity can be used to serve a class of traffic called scavenger class [12], [13], which is under-prioritized compared to even the best-effort traffic. The scavenger traffic gets little delay assurance, and is served using the spare capacity. We analyze the pricing of such traffic as a way to mitigate the impact of time-constraint restriction on access price.

Let S be a set of scavenger class flows served in addition to the best-effort flows \mathcal{F} by the ISP. Let the scavenger flows have a utility $\sigma^S(1-\beta)^{-1}z^{1-\beta}$ as a function of the data-volume z allocated to each flow in the class. A natural comparison of revenue from scavenger class R_S is with the additional revenue $R_u^* - R_p^*$ that the ISP would have earned from best-effort flows $f \in \mathcal{F}$ with time-flexible pricing, allowing full utilization of capacity at all times. A sufficiently high scavenger class flow utility level σ^S allows the ISP to regain the revenue loss due to time-constrained pricing as shown by the following theorem, proved in Appendix A.

Theorem 4: Assume utility functions associated with best-effort flows are alpha-fair (4) with $\alpha_f = \alpha$ for all $f \in \mathcal{F}$ and flow-independent utility levels $\{\sigma^t\}$. Assume utility functions associated with scavenger flows are alpha-fair (4) with $\alpha_f = \beta$ for all $f \in S$ and flow-independent utility level σ^S . Then, revenue from scavenger class R_S is at least as much as the revenue gain from time-flexible pricing $R_u^* - R_p^*$ if

$$\sigma^S \geq \sigma^m \frac{\|\mathcal{F}\|^\alpha}{\|\mathcal{S}\|^\beta} \frac{1-\beta}{1-\alpha} C^{\beta-\alpha} \frac{\sum_t \left(\frac{\sigma^t}{\sigma^m} - \left(\frac{\sigma^t}{\sigma^m}\right)^{1/\alpha}\right)}{\left(\sum_t \left(1 - \left(\frac{\sigma^t}{\sigma^m}\right)^{1/\alpha}\right)\right)^{1-\beta}} \quad (15)$$

C. An Illustrative Example

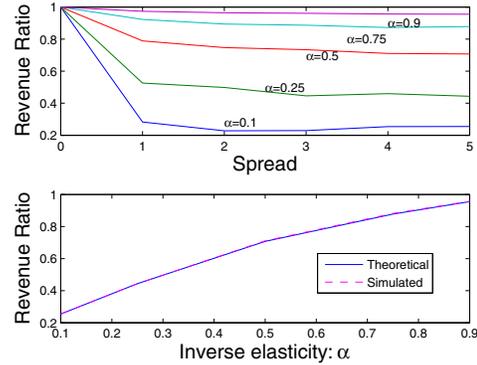


Fig. 2. Comparison of revenue loss from time constrained pricing at different spread in utility levels and different consumer demand elasticity.

Consider the example in Section III-B. The upper part of Figure 2 demonstrates the revenue loss from time-constrained pricing by plotting the revenue ratio R_p^*/R_u^* against the spread in utility levels, represented by $\sigma_1 - \sigma_0$. Utility levels were realized through a uniform distribution for 10 flows. Higher spread results in higher reduction in revenue, with the loss saturating at higher values of the spread. For σ with a uniform distribution in $[\sigma_0, \sigma_1]$, we have

$$E(\sigma^{(1/\alpha)}) = \frac{\sigma_1^{(1+1/\alpha)} - \sigma_0^{(1+1/\alpha)}}{1 + 1/\alpha}, \quad E(\sigma) = \frac{\sigma_1^2 - \sigma_0^2}{2}$$

At high values of $\sigma_1 - \sigma_0$, the revenue ratio can be shown to saturate to

$$\frac{R_p^*}{R_u^*} = \frac{\sigma_1 E(\sigma^{(1/\alpha)})}{\sigma_1^{(1/\alpha)} E(\sigma)} \rightarrow \frac{2}{1 + 1/\alpha}$$

The revenue loss is plotted against α in the lower part of Figure 2. Low values of α , indicating high price sensitivity, results in high losses.

V. NON-DISCRIMINATORY PRICING ACROSS FLOWS

In Section IV, we considered restriction over *time* on the prices r_f^t charged by the monopolist ISP. However, the ISP could discriminate across *flows*, practicing first-order price discrimination [20] to extract the maximum revenue. In practice, the ability of a monopoly ISP to practice such complete price discrimination across flows is restricted. The possibility of regulatory scrutiny and demand for “network neutrality” [6] prevents monopoly ISPs from practicing extensive price discrimination. We investigate the revenue loss from such price restriction across flows and mitigation of the loss through non-linear pricing.

A. Revenue Loss from Uniform Pricing across Flows

The revenue loss from price restriction across flows is incurred due to variations in the consumer utility functions. We simplify the model by removing the time-dependency of the utility levels. As explained in section II, we consider a probabilistic model of utility levels across consumers, governed by the density distribution function $f(\sigma)$ defined over a range $[\sigma_0, \sigma_1]$ with the corresponding cumulative distribution function $F(\sigma) = \int_{\sigma_0}^{\sigma} f(\sigma)$. With no price restriction, let the price $r(\sigma) = g(\sigma) + h(\sigma)x(\sigma)$ be the σ dependent price charged to a flow with utility level σ . The demand function of the consumer facing this price is given by $y(\sigma) = u'^{-1}(h(\sigma)/\sigma)$. The revenue maximization problem for the ISP in this case is similar to (6), with the ISP capturing the entire consumer net-utility by setting the flat fee $g(\sigma) = \sigma u(y(\sigma)) - h(\sigma)y(\sigma)$. The revenue from unrestricted pricing, which we denote by Π_u , is then determined by the appropriate usage fee, and is given by

$$\begin{aligned} & \text{maximize} && \int_{\sigma_0}^{\sigma_1} \sigma u(y(\sigma)) f(\sigma) d\sigma \\ & \text{subject to} && \int_{\sigma_0}^{\sigma_1} y(\sigma) f(\sigma) d\sigma \leq C \\ & \text{variables} && \{h(\sigma)\} \end{aligned} \quad (16)$$

For alpha-fair utility functions (4), the demand function is given by

$$y(\sigma) = (\sigma/h(\sigma))^{1/\alpha} \quad (17)$$

The following theorem states the solution to problem (16) in this case.

Theorem 5: Assuming utility functions are alpha-fair (4) with $\alpha_f = \alpha$ for the continuum of flows, the solution to problem (16) is given by

$$\begin{aligned} h^{1/\alpha} &= \frac{B(\sigma_0, \sigma_1)}{C} \\ B(\sigma_0, \sigma_1) &= \int_{\sigma_0}^{\sigma_1} (\sigma)^{1/\alpha} f(\sigma) d\sigma \\ \Pi_u &= (1 - \alpha)^{-1} h^{1-1/\alpha} B(\sigma_0, \sigma_1) = (1 - \alpha)^{-1} hC \end{aligned} \quad (18)$$

The above result can be readily shown: The optimal price structure in (18) represents the KKT conditions [21] for the optimization problem in (16).

With the ISP restricted to charge a uniform price with linear increase in usage, the flat component g and the usage rate

h do not change with the consumer utility level σ . The σ -independent usage and flat rate pricing does not allow the ISP to capture the consumer net-utility at every utility level, resulting in a revenue loss. The net-utility of a consumer with utility level σ given by $\sigma u(y(\sigma)) - h y(\sigma)$ can be shown to be increasing in σ . Given the restriction that the user net-utility cannot be negative, the ISP has to set the σ -independent flat-rate component g to the net-utility of the consumer with the minimum utility level $\sigma_T \in [\sigma_0, \sigma_1]$ that the ISP is willing to serve. Consumers with utility level $\sigma < \sigma_T$ are not served by the ISP. The threshold utility level σ_T is then part of the revenue optimization problem, with higher σ_T providing a higher flat-fee component from all consumers with utility levels higher than σ_T , but also resulting in zero revenue from all consumers with utility level less than σ_T . The revenue maximization problem with this price restriction, whose optimum we denote by Π_r , is given by

$$\begin{aligned} & \text{maximize} && \int_{\sigma_T}^{\sigma_1} (g + h y(\sigma)) f(\sigma) d\sigma \\ & \text{subject to} && \int_{\sigma_T}^{\sigma_1} y(\sigma) f(\sigma) d\sigma \leq C, \\ & && g = \sigma_T u(y(\sigma_T)) - h y(\sigma_T) \\ & \text{variables} && \{\sigma_T, h\} \end{aligned} \quad (19)$$

For alpha-fair utility functions of the form in (4), the marginal utility at zero data allocation is infinity, resulting in infinite penalty for not serving a flow of non-zero utility level. Indeed, we show in Appendix B that $\sigma_T = \sigma_0$, so that the ISP is better off serving all flows, and further prove that the solution to (19) is given by the following theorem.

Theorem 6: Assuming alpha-fair utility functions of the form in (4) with $\alpha_f = \alpha$ for the continuum of flows, the revenue, from uniform price across flows with linear increase in usage, is given by

$$\begin{aligned} \Pi_r &= hC \left(1 + \frac{\alpha \sigma_0^{1/\alpha}}{(1-\alpha)B(\sigma_0, \sigma_1)} \right) \\ h^{1/\alpha} &= \frac{B(\sigma_0, \sigma_1)}{C} \\ B(\sigma_0, \sigma_1) &= \int_{\sigma_0}^{\sigma_1} (\sigma)^{1/\alpha} f(\sigma) d\sigma \end{aligned} \quad (20)$$

and the ratio of revenue from flow-restricted pricing Π_r to revenue from unrestricted pricing Π_u is given by

$$\frac{\Pi_r}{\Pi_u} = 1 - \alpha \left(1 - \frac{\sigma_0^{1/\alpha}}{B(\sigma_0, \sigma_1)} \right) \quad (21)$$

Remarks: Notice that $\sigma_1^{1/\alpha} \geq B(\sigma_0, \sigma_1) \geq \sigma_0^{1/\alpha}$. Therefore, $\Pi_r \leq \Pi_u$ with equality when all consumers have the same utility level σ_0 ($\sigma_1 = \sigma_0$). The revenue inefficiency stems from the inability of the ISP to charge the flow-dependent flat fee to capture the net-utility of each flow. The inefficiency is higher with higher values of α since a larger fraction of the revenue is dependent on the flat fee with higher α , as shown in Theorem 2.

The impact of the spread in utility levels on the revenue can be estimated by noting the following lower bound on the revenue inefficiency.

$$\frac{\Pi_r}{\Pi_u} \geq 1 - \alpha \left(1 - \frac{\sigma_0^{1/\alpha}}{\sigma_1^{1/\alpha}} \right) \quad (22)$$

The lower bound can be interpreted as the inverse competitive ratio of the restricted pricing scheme. The revenue inefficiency is low if the spread in utility levels is low. A large spread in

utility levels among flows can, however, result in high revenue inefficiencies.

B. Non-linear Pricing for Revenue Loss Mitigation

The price restriction across flows has thus far been modeled as a combination of flat-fee and usage fee, resulting in a price linear in usage $r(\sigma) = g + hy(\sigma)$. The ISP can incur a higher revenue by offering a price that is non-linear in usage *without* discriminating on a per flow basis. The non-linear price typically takes the form of quantity discounts where the price per unit data rate is discounted for higher data rate purchases. We study the revenue outcome of non-linear pricing in this section by invoking results from second-order price discrimination in economics [20].

The non-linear price offering is in the form of a “package” consisting of the data rate $y(\sigma)$ and the price $r(\sigma)$ offered to the consumers. The package has to be structured in such a way that the package $\{y(\sigma), r(\sigma)\}$ for a given σ , intended for a consumer with utility level σ , should indeed be the most desirable choice for the consumer. This requirement is referred to as the incentive compatibility and the requirement can be expressed in terms of the net-utility of consumers as

$$\sigma u(y(\sigma)) - r(\sigma) \geq \sigma u(y(\tilde{\sigma})) - r(\tilde{\sigma}), \forall \tilde{\sigma} \in [\sigma_0, \sigma_1] \quad (23)$$

In addition, the package has to be structured so that the consumers derive non-negative net-utility from the package, a requirement that is referred to as the individual rationality, expressed in terms of net-utility of consumers as

$$\sigma u(y(\sigma)) - r(\sigma) \geq 0, \forall \sigma \in [\sigma_0, \sigma_1] \quad (24)$$

The revenue maximization problem for the ISP, whose maximum is denoted by Π_n , can then be written as

$$\begin{aligned} & \text{maximize} && \int_{\sigma_0}^{\sigma_1} r(\sigma) f(\sigma) d\sigma \\ & \text{subject to} && \int_{\sigma_0}^{\sigma_1} y(\sigma) f(\sigma) d\sigma \leq C \\ & && \sigma u(y(\sigma)) - r(\sigma) \geq \sigma u(y(\tilde{\sigma})) - r(\tilde{\sigma}), \forall \tilde{\sigma} \in [\sigma_0, \sigma_1] \\ & && \sigma u(y(\sigma)) - r(\sigma) \geq 0, \\ & \text{variables} && \{y(\sigma), r(\sigma)\} \forall \sigma \in [\sigma_0, \sigma_1] \end{aligned} \quad (25)$$

In general, we expect $\Pi_u \geq \Pi_n \geq \Pi_r$. Consider a special case when consumer utility function is uniformly distributed in the range $[0, 1]$ so that

$$f(\sigma) = 1, \quad F(\sigma) = 1 - \sigma, \quad \sigma_0 = 0, \quad \sigma_1 = 1 \quad (26)$$

For this special case, we can state the following result, proved in Appendix C, on the revenue ratios.

Theorem 7: Assuming utility functions are alpha-fair (4) with $\alpha_f = \alpha$ for the continuum of flows, and the utility levels distributed uniformly as in (26), we have

$$\frac{\Pi_r}{\Pi_n} = 2^\alpha(1 - \alpha), \quad \frac{\Pi_r}{\Pi_u} = (1 - \alpha), \quad \frac{\Pi_n}{\Pi_u} = 2^{-\alpha} \quad (27)$$

C. An Illustrative Example

Consider the example in Section III-B. The upper part of Figure 3 plots the revenue from flow-restricted pricing, unrestricted pricing, and non-linear pricing. The lower part plots the revenue ratio between restricted and non-linear pricing, restricted and unrestricted pricing, and non-linear and unrestricted

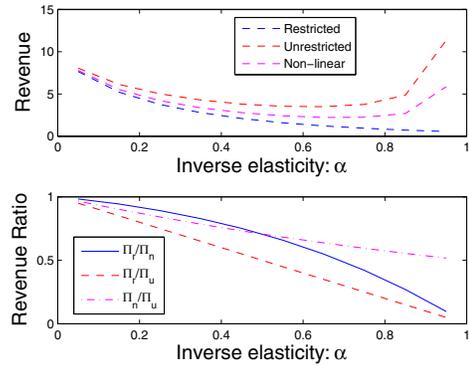


Fig. 3. Comparison of the ISP revenue from three cases: unrestricted, restricted and non-linear pricing.

pricing. Non-linear pricing allows the ISP to recover losses from restrictions on price differentiation between flows, with the loss no worse than 50%.

VI. CONCLUSION

The questions of “whom to price?” and “how to price?” are two related aspects of connectivity economics, that has bearing on the expansion of broadband access for ubiquitous availability of Internet connectivity and preservation of open Internet. The first question is addressed in previous works, including our recent one [19], in the context of two-sided pricing by investigating the appropriate price split between end-users and content (or application) providers. This paper addressed the second question in the context of a monopoly ISP with complete pricing power, charging one-sided access price to consumers of connectivity that can be either the end-users or the content providers.

The work can be extended for further analysis of benefits distribution under market scenarios where ISPs have less pricing power, consumers are price anticipating [15], or ISPs have insufficient information on consumer utilities [11]. Further analysis is required to incorporate inter-ISP interactions [16] when consumer data packets are routed through multiple bottleneck links operated by different ISPs.

A different line of work is the analysis of access price when the ISP offers tiered connectivity services with quality of service differentiation across tiers. Related works have analyzed the need for service tiering [2] or analyze the revenue gain from a specific form of service tiering [17], [18]. A topic of future interest is the optimization of key parameters of service differentiation *jointly*, with pricing.

Although the analysis in this paper is applicable to wireless networks with the fixed capacity representing the average shared data-rate, further refinement is required to analyze the effect of interference and wireless channel variations across consumers and over time. A related line of work is the analysis of pricing schemes for secondary spectrum access through cognitive radio.

Beyond revenue maximization, access prices are ultimately determined by additional factors, including the need to recover

the substantial capital cost of provisioning the access link. A clearer picture on access pricing will emerge with further analysis that takes the ISP cost model into account.

APPENDIX A: PROOFS OF THEOREMS 3 AND 4

The first-order KKT conditions for problem (11) result in

$$h_f = \frac{\sum_t \mu^t \frac{\partial x_f^t}{\partial h_f}}{\sum_t \frac{\partial x_f^t}{\partial h_f}} = \frac{\sum_t \mu^t \eta_f^t x_f^t}{\sum_t \eta_f^t x_f^t} \quad (28)$$

where μ^t is the dual congestion price associated with the capacity constraint at each instant t and η_f^t is the elasticity of demand (3). With utility functions specialized to the form in (4), we have

$$h_f = \frac{\sum_t \mu^t (\sigma_f^t)^{1/\alpha_f}}{\sum_t (\sigma_f^t)^{1/\alpha_f}} \quad (29)$$

The complementary slackness condition implies that μ^t is such that

$$\begin{aligned} \mu^t &= 0 & \text{if } \sum_f (\sigma_f^t/h_f)^{1/\alpha_f} < C \\ \mu^t &> 0 & \text{if } \sum_f (\sigma_f^t/h_f)^{1/\alpha_f} = C \end{aligned} \quad (30)$$

Let \mathcal{T} be the subset of times with peak aggregate usage, so that we have $\sum_f (\sigma_f^t/h_f)^{1/\alpha_f} = C, t \in \mathcal{T}$. It follows that

$$h_f = \frac{\sum_{t \in \mathcal{T}} \mu^t (\sigma_f^t)^{1/\alpha_f}}{\sum_t (\sigma_f^t)^{1/\alpha_f}}$$

To evaluate the effect of the time factors σ_f^t , consider flow-independent time factors σ^t , and flow independent utility function of the form in (4) parameterized by α . Flows are identical so that the capacity share is identical. With $\sigma^m = \max\{\sigma^t\}$ the maximum utility level for any flow over time, we have $\sigma^t = \sigma^m, \forall t \in \mathcal{T}$. The time-constrained price h_f will have to be set so as to limit the data rates at these times to be within capacity, resulting in

$$h = \sigma^m (C/\|\mathcal{F}\|)^{-\alpha} \quad (31)$$

where $\|\mathcal{F}\|$ is the number of flows sharing the capacity. The data rate for each flow at any time t is then given by

$$x^t = \left(\frac{\sigma^t}{\sigma^m} \right)^{1/\alpha} (C/\|\mathcal{F}\|) \quad (32)$$

With no time-constraint on pricing, the usage price at each time instant h^t can be set so that the data rate for each flow is

$$x = (C/\|\mathcal{F}\|) \quad (33)$$

The ratio of revenue R_p^* to the revenue R_u^* without time-constrained pricing is then given by

$$\begin{aligned} \frac{R_p^*}{R_u^*} &= \frac{\sum_t \sigma^t u(x^t)}{\sum_t \sigma^t u(x)} = \frac{\sum_t \sigma^t (x^t)^{1-\alpha}}{\sum_t \sigma^t (x)^{1-\alpha}} \\ &= \frac{\sum_t \sigma^t \left(\frac{\sigma^t}{\sigma^m} \right)^{1/\alpha(1-\alpha)} (C/\|\mathcal{F}\|)^{1-\alpha}}{\sum_t \sigma^t ((C/\|\mathcal{F}\|))^{1-\alpha}} \\ &= \frac{\sum_t \sigma^t \left(\frac{\sigma^t}{\sigma^m} \right)^{1/\alpha(1-\alpha)}}{\sum_t \sigma^t} \end{aligned} \quad (34)$$

yielding the result in (12).

For proof of Theorem 4, notice that data rate allocation for best-effort traffic flows is given by (32). With the spare capacity from time-constrained price equally allocated across scavenger class traffic, the capacity share of each traffic flow is given by

$$z = \frac{\sum_t (C - \sum_f x_f^t)}{\|\mathcal{S}\|} = \frac{C}{\|\mathcal{S}\|} \sum_t \left(1 - \left(\frac{\sigma^t}{\sigma^m} \right)^{1/\alpha} \right) \quad (35)$$

A combination of flat and usage fee allows the monopoly ISP to extract all the net-utility from the scavenger class users, generating a revenue R_S of

$$\|\mathcal{S}\| \sigma^S (1-\beta)^{-1} \left(\frac{C}{\|\mathcal{S}\|} \right)^{1-\beta} \left(\sum_t \left(1 - \left(\frac{\sigma^t}{\sigma^m} \right)^{1/\alpha} \right) \right)^{1-\beta} \quad (36)$$

The revenue $R_u^* - R_p^*$ is given by

$$\|\mathcal{F}\| \sigma^m (1-\alpha)^{-1} \left(\frac{C}{\|\mathcal{F}\|} \right)^{1-\alpha} \sum_t \left(\frac{\sigma^t}{\sigma^m} - \left(\frac{\sigma^t}{\sigma^m} \right)^{1/\alpha} \right) \quad (37)$$

from which the result in (15) follows.

APPENDIX B: PROOF OF THEOREM 6

Consider the Lagrangian for the problem (19)

$$\begin{aligned} \mathcal{L}(h, \sigma_T, \mu) &= \int_{\sigma_T}^{\sigma_1} (\sigma_T u(y(\sigma_T)) - h y(\sigma_T) + h y(\sigma)) f(\sigma) d\sigma \\ &+ \mu (C - \int_{\sigma_T}^{\sigma_1} y(\sigma) f(\sigma) d\sigma) \\ &= \int_{\sigma_T}^{\sigma_1} \left(\sigma_T (1-\alpha)^{-1} \left(\frac{\sigma_T}{h} \right)^{1/\alpha-1} - h \left(\frac{\sigma_T}{h} \right)^{1/\alpha} + h \left(\frac{\sigma}{h} \right)^{1/\alpha} \right) f(\sigma) d\sigma \\ &+ \mu \left(C - \int_{\sigma_T}^{\sigma_1} \left(\frac{\sigma}{h} \right)^{1/\alpha} f(\sigma) d\sigma \right) \end{aligned}$$

The usage price h should be such that $\frac{\partial \mathcal{L}}{\partial h} = 0$ which translates to

$$(\mu h^{-1} - 1 + \alpha) \int_{\sigma_T}^{\sigma_1} \sigma^{1/\alpha} f(\sigma) d\sigma = \alpha \sigma_T^{1/\alpha} (1 - F(\sigma_T)) \quad (38)$$

Since $\sigma_T \leq \sigma_1$, we have $(\mu h^{-1} - 1 + \alpha) \leq \alpha$ or $h \geq \mu$.

The threshold utility level σ_T should be such that $\frac{\partial \mathcal{L}}{\partial \sigma_T} \geq 0$. We have

$$\frac{\partial \mathcal{L}}{\partial \sigma_T} = f(\sigma_T) \left(-\sigma_T (1-\alpha)^{-1} \left(\frac{\sigma_T}{h} \right)^{1/\alpha-1} + \mu \left(\frac{\sigma_T}{h} \right)^{1/\alpha} \right) \quad (39)$$

which results in the requirement that

$$h \leq \mu(1-\alpha) \quad (40)$$

For $0 < \alpha < 1$, this contradicts the requirement that the optimum usage fee be no less than the congestion price: $h \geq \mu$. That is, there is no benefit to setting the threshold utility level σ_T to any value higher than σ_0 , since the loss of revenue from not serving a consumer of utility level σ_T is more than the gain in revenue from distributing the resulting spare capacity to higher utility level consumers.

Setting $\sigma_T = \sigma_0$, the revenue from price restriction across flows, Π_r , is given by

$$\int_{\sigma_0}^{\sigma_1} \left(\sigma_0 (1-\alpha)^{-1} \left(\frac{\sigma_0}{h} \right)^{1/\alpha-1} - h \left(\frac{\sigma_0}{h} \right)^{1/\alpha} + h \left(\frac{\sigma}{h} \right)^{1/\alpha} \right) f(\sigma) d\sigma \quad (41)$$

with the usage price h set to satisfy the capacity constraint. The result in Theorem 6 follows.

APPENDIX C: PROOF OF THEOREM 7

We first invoke a result on second-degree price discrimination [14] from economics.

Lemma 1: Under certain technical conditions on consumer utility function $v(x, \sigma)$, the data rate allocation $y(\sigma)$ under second-degree price discrimination satisfies

$$\frac{\partial v(y(\sigma), \sigma)}{\partial x} - \left(\mu + \frac{1 - F(\sigma)}{f(\sigma)} \frac{\partial^2 v(y(\sigma), \sigma)}{\partial x \partial \sigma} \right) = 0 \quad (42)$$

where μ is the capacity price. Further, the form of the price structure $r(\sigma)$ is given by

$$r(\sigma) = v(y(\sigma), \sigma) - \int_{\sigma_0}^{\sigma} \frac{\partial v(y(t), t)}{\partial \sigma} dt \quad (43)$$

For $v(x, \sigma) = \sigma u(x)$ where $u(x)$ is the alpha-fair utility function as defined in (4) with $\alpha_f = \alpha$, we have

$$\begin{aligned} \frac{\partial v(x, \sigma)}{\partial x} &= \sigma x^{-\alpha} \\ \frac{\partial^2 v(y(\sigma), \sigma)}{\partial x \partial \sigma} &= x^{-\alpha} \end{aligned} \quad (44)$$

Plugging in (42), we get

$$y(\sigma) = \left(\mu^{-1} \left(\sigma - \frac{1 - F(\sigma)}{f(\sigma)} \right) \right)^{1/\alpha} \quad (45)$$

where μ is the congestion price such that

$$\int_{\sigma_0}^{\sigma_1} y(\sigma) d\sigma = C \quad (46)$$

We see that in comparison to demand in the linear pricing case (17), the demand in the non-linear pricing case is distorted downward for every utility level σ except the highest $\sigma = \sigma_1$, at which point we have $1 - F(\sigma_1) = 0$. The relative share of capacity has been adjusted in favor of consumers with higher utility levels.

Consider the case when consumer utility function is uniformly distributed in the range $[0, 1]$ with resulting distribution functions give by (26). In this case, we have

$$y(\sigma) = \begin{cases} \left(\frac{2\sigma-1}{\mu} \right)^{1/\alpha} & \text{if } \sigma \in [0.5, 1] \\ 0 & \text{otherwise} \end{cases} \quad (47)$$

where

$$\mu = (2C(1 + 1/\alpha))^{-\alpha} \quad (48)$$

We see that consumers with utility level $\sigma \in [0, 0.5]$ are not being served by the ISP charging a non-linear access price. Plugging (47) into (43), the price structure $r(\sigma)$ for $\sigma \in [0.5, 1]$ can be easily shown to be given by

$$r(\sigma) = \frac{(2C(1 + 1/\alpha))^{1-\alpha}}{(1-\alpha)} (2\sigma-1)^{1/\alpha} \left(\frac{\sigma}{2\sigma-1} - \alpha/2 \right) \quad (49)$$

The resulting revenue from non-linear pricing is given by

$$\Pi_n = \left(\frac{\alpha}{2C(1 + \alpha)} \right)^\alpha \frac{C}{1-\alpha} \quad (50)$$

Plugging in parameters from (26) of the assumed uniform distribution into (20), we obtain the revenue from restricted pricing as

$$\Pi_r = \left(\frac{\alpha}{C(1 + \alpha)} \right)^\alpha C \quad (51)$$

Similarly, plugging parameters from (26) of the assumed uniform distribution into (18), we obtain the revenue from unrestricted pricing as

$$\Pi_u = \left(\frac{\alpha}{C(1 + \alpha)} \right)^\alpha \frac{C}{1-\alpha} \quad (52)$$

thereby concluding the proof.

ACKNOWLEDGMENT

The authors would like to acknowledge helpful discussions with Sundeep Rangan at Qualcomm Inc. This work has been in part supported by NSF CNS-0905086 and CNS-0519880.

REFERENCES

- [1] J. K. Mackie-Mason and H. Varian, "Pricing Congestible Network Resources," *IEEE Journal on Selected Areas in Communication*, vol. 13, no. 7, pp. 1141-1149, 1995.
- [2] G. Kesidis, A. Das and G. de Veciana, "On Flat-rate and Usage-based Pricing for Tiered Commodity Internet Services", *Proc. CISS*, March 2008.
- [3] A. M. Odlyzko, "Internet pricing and the history of communications", *Computer Networks* vol. 36, pp. 493-517, 2001.
- [4] FCC, "Open Internet Notice of Proposed Rulemaking", <http://www.openinternet.gov/about-the-nprm.html>, October 2009.
- [5] CRTC, "Review of the Internet traffic management practices of Internet service providers", <http://www.crtc.gc.ca/eng/archive/2009/2009-657.htm>, October 2009.
- [6] C. Yoo, "Network Neutrality and the Economics of Congestion", *Georgetown Law Journal*, Vol. 94, June 2006.
- [7] M. Falkner, M. Devetsikiotis, and I. Lambadaris, "An overview of pricing concepts for broadband IP networks", *IEEE Communications Surveys*, vol. 3, no. 2, 2000.
- [8] W. Oi, "A Disneyland dilemma: two-part tariffs for a Mickey Mouse monopoly", *Quarterly Journal of Economics*, 1971.
- [9] S. Scotchmer, "Profit-maximizing clubs", *Journal of Public Economics*, 27, 2545, 1985.
- [10] T. Basar, R. Srikant, "Revenue-Maximizing Pricing and Capacity Expansion in a Many-Users Regime", *Proc. of IEEE Infocom*, 2002.
- [11] L. Jiang, S. Parekh, J. Walrand, "Time-dependent Network Pricing and Bandwidth Trading", *IEEE International Workshop on Bandwidth on Demand*, 2008.
- [12] Internet 2, "QBone Scavenger Service (QBSS)", <http://qbone.internet2.edu/qbss/>
- [13] R. Bless, K. Wehrle, "A Lower Than Best-Effort Per-Hop Behavior", Internet draft at <http://tools.ietf.org/html/draft-bless-diffserv-lbe-phb-00>, September 1999.
- [14] E. Maskin, J. Riley, "Monopoly with Incomplete Information", *RAND Journal of Economics*, Vol. 15 (1984), pp. 171-196.
- [15] R. Johari, J. N. Tsitsiklis, "Efficiency of Scalar-Parameterized Mechanisms", *Operations Research*, Vol. 57, No. 4, pp. 823-839, July-August 2009.
- [16] R. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, D. Rubenstein, "Interconnecting Eyeballs to Content: A Shapley Value Perspective on ISP Peering and Settlement", *Proc. ACM NetEcon*, 2008.
- [17] D. Acemoglu, A. Ozdaglar, and R. Srikant, "The marginal user principle for resource allocation in wireless networks," *Proc. 43rd IEEE Conference on Decision and Control* December 2004.
- [18] S. Shakkottai, R. Srikant, A. Ozdaglar and D. Acemoglu, "The Price of Simplicity", *IEEE Journal on Selected Areas in Communication*, Vol. 26, No. 7, September 2008.
- [19] P. Hande, M. Chiang, R. Calderbank, S. Rangan, "Network Pricing and Rate Allocation with Content Provider Participation", *Proc. of IEEE Infocom*, April 2009.
- [20] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*, Oxford University Press, Oxford, United Kingdom, 1995.
- [21] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.