

An Economic Model for Pricing Tiered Network Services*

Qian Lv, George N. Rouskas
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-8206, USA
E-mail: {qlv,rouskas}@ncsu.edu

September 30, 2009

Abstract

We consider networks offering tiered services and corresponding price structures, a model that has become prevalent in practice. We develop an economic model for such networks and make contributions in two important areas. First, we formulate the problem of selecting the service tiers from three perspectives: one that considers the users' interests only, one that considers only the service provider's interests, and one that considers both simultaneously, i.e., the interests of society as a whole. We also present an approximate yet accurate and efficient solution approach for tackling these nonlinear programming problems. Given the set of (near-) optimal service tiers, we then employ game-theoretic techniques to find an optimal price for each service tier that strikes a balance between the conflicting objectives of users and service provider. This work provides a theoretical framework for reasoning about and pricing Internet tiered services, as well as a practical toolset for network providers to develop customized menus of service offerings. Our results also indicate that tiering solutions currently adopted by ISPs perform poorly both for the providers and the society overall.

1 Introduction

Internet service providers (ISPs) have introduced several forms of a *tiered service*, in which users may select from a small set of tiers which offer progressively higher levels of service with a corresponding increase in price. Multitiered price systems are prevalent for both business and residential Internet access, and have been employed in various forms regardless of whether the underlying pricing scheme is capacity-based (in which the subscription fee is determined solely by the user's access speed) or usage-sensitive (in which price is a function of the actual bytes transferred over a certain time period, usually one month).

If designed and applied appropriately, multitiered pricing schemes have the potential to be a catalyst for Internet service innovation and penetration. On the provider side, tiered structures can be an effective tool for ISPs to optimize and specialize their offerings so as to capitalize on the increasing sophistication and requirements of various segments of Internet users, as well as to differentiate themselves from the competition. On the user side, tiered pricing is likely to spur adoption by providing a wide menu of customized services from which users may select based on needs and affordability.

To realize this potential, it is crucial that both the service tiers and the corresponding prices be determined in a manner that takes into account simultaneously the (usually conflicting) objectives of users and providers.

*This work was supported by the NSF under grant CNS-0434975.

In current practice, however, there is considerable lack of transparency in how ISPs set their tiered price structures, and it is unclear whether the perspective of users is even considered in the process. For instance, certain service tiers for business Internet access are based on the bandwidth hierarchy of the underlying network infrastructure (e.g., DS-1, OC-3, etc.). While this is a natural arrangement for the service provider, it is unlikely that hierarchical rates designed decades ago for voice traffic would be a good match for today’s business data applications. The ADSL tiers available from various providers, on the other hand, appear to have been determined in some *ad-hoc* manner; similar observations apply to the usage-sensitive tiered price structure employed by one cable ISP for a recent pilot program [6], which sets the tiers at 5 GB, 10 GB, 20 GB, and 40 GB of monthly download traffic. While tier values that are round numbers may be an appropriate choice for marketing purposes, the relationship between these exponentially increasing levels of service (and their price) and the usage patterns (and willingness or ability to pay) of the population of potential subscribers is open to debate.

In this work, we develop an economic model for networks with a tiered service and price structure and make contributions in two important areas. First, we present nonlinear programming formulations for the problem of selecting the service tiers from three perspectives: one that considers the users’ interests only, one that considers only the service provider’s interests, and one that considers both simultaneously, i.e., the interests of society as a whole. We also devise an efficient solution methodology by developing a dynamic programming algorithm to solve optimally an approximate formulation of the original service tier selection problem. Given the set of (near-) optimal service tiers, we then employ game-theoretic techniques based on Nash bargaining to find an optimal price for each service tier that strikes a balance between the conflicting objectives of users and service provider. Our work provides a theoretical framework for reasoning about and pricing Internet tiered services, as well as a practical toolset for network providers to develop customized menus of service offerings that cater to user needs while ensuring that both parties are satisfied.

The rest of the paper is organized as follows. In Section 2, we describe the tiered-service network we consider in this study, along with related applications. In Section 3, we introduce an economic model for tiered-service networks that takes into account the user’s perspective, the provider’s perspective, or both. We also formulate and solve the corresponding problems for selecting the set of service tiers optimally. In Section 4 we use Nash bargaining theory to determine an optimal price for each of the service tiers. We present numerical results in Section 5, and we point to extensions of our work in Section 6. We discuss related work in Section 7, and we conclude the paper in Section 8.

2 Tiered-Service Networks

We consider a network that offers a service characterized by a single parameter, e.g., the bandwidth of the user’s access link or the amount of traffic generated by the user, and charges users on the basis of the amount of service they receive. Users may request any amount of service depending on their needs and their willingness or ability to pay the corresponding service fee. We assume that the distribution of the size x of user service requests is known; such a distribution may be obtained empirically, or extrapolated from observed user behavior and application requirements. Let $f(x)$ and $F(x)$ be the probability density function (pdf) and cumulative distribution function (cdf), respectively, representing the population of user requests. The pdf and cdf are defined in the interval $[x_{min}, x_{max}]$, where x_{min} and x_{max} correspond to the minimum and maximum, respectively, amount of service requested by any user.

The network offers K levels (tiers) of service, where typically K is a small integer, much smaller than the number N of network users (i.e., $K \ll N$). We define $Z = \langle z_1, z_2, \dots, z_K \rangle$ as the vector of service tiers offered by the network provider; without loss of generality, we assume that the service tiers are distinct and are labeled such that $z_1 < z_2 < \dots < z_K$. For notational convenience, we also define the “null” service tier $z_0 = 0$.

With tiered service, a user with service request x , $x_{min} \leq x \leq x_{max}$, subscribes to service tier z_j such




Figure 1: Sample mapping of service requests to service tiers

that $z_{j-1} < x \leq z_j$. Figure 1 shows a sample mapping of service requests to a vector of $K = 6$ service tiers under which all users with requests $x \in (z_{j-1}, z_j]$ subscribe to service tier z_j , $j = 1, \dots, K$. Note also that in order to accommodate all user requests, the highest tier must be such that $z_K = x_{max}$, an assumption we will make throughout this work.

The network provider incurs a cost for the service it provides, and consequently, it will be inclined to select the service tiers, and the corresponding price to charge, so as to recoup its costs (and make a profit). On the other hand, each user subscribes to a service that is at least as good as the one requested, but the additional value, if any, that the user receives may be offset by the higher cost of the service. Our aim is to apply economic theory to capture analytically these tradeoffs, and to develop techniques to select the service tiers and prices in a manner that accounts for both the users' and providers' perspectives.

To develop an economic model for tiered-service networks, we assume the existence of three non-decreasing functions of service x , as shown in Figure 2. The *utility* function, $U(x)$, is a measure of the value that users receive from the service, and it stands for their willingness to pay for the service. The *cost* function, $C(x)$, represents the cost incurred by the provider for offering the service. Finally, the *price* function, $P(x)$, represents the amount that the service provider charges for the service. Figure 2 shows that $U(x)$ lies above $P(x)$ (otherwise users would not be willing to pay for the service), and in turn $P(x)$ lies above $C(x)$ (otherwise providers would not be inclined to offer the service). We make the reasonable assumption that utility, cost, and price are all expressed in the same units (e.g., US\$). Note that utility and cost typically depend only on the user and service provider, respectively, but that price is the result of market dynamics and the relative bargaining power of users and service providers.

In our model, we make two fundamental assumptions. First, we consider a *homogeneous* user population in the sense that all users are characterized by the same utility function $U(x)$; the case of heterogeneous users is the subject of ongoing research in our group. Second, we assume that any user with demand x that is mapped to some tier z_j will subscribe to this tier as long as the price is less than its willingness to pay as represented by the utility value $U(z_j)$.

The tiered service model we consider in this paper arises naturally under both pricing schemes, capacity-based or usage-sensitive, that are prevalent for Internet services [8].

Capacity-Based Pricing. Capacity-based schemes relate pricing to usage by setting a price based on the bandwidth or speed of the user's connection link. This is accomplished by charging for the configuration (i.e., bandwidth) of the connection, but not the actual bits sent or received. Capacity-based pricing is the prevailing pricing policy for residential broadband Internet access services. This scheme relates to our tiered service model as follows: the service is characterized by the amount of access bandwidth, each of the service tiers z_1, \dots, z_K , corresponds to a certain access speed, and users are charged based on the tier to which they have subscribed. Currently, the offered tiers are either tied to the bandwidth hierarchy of the underlying

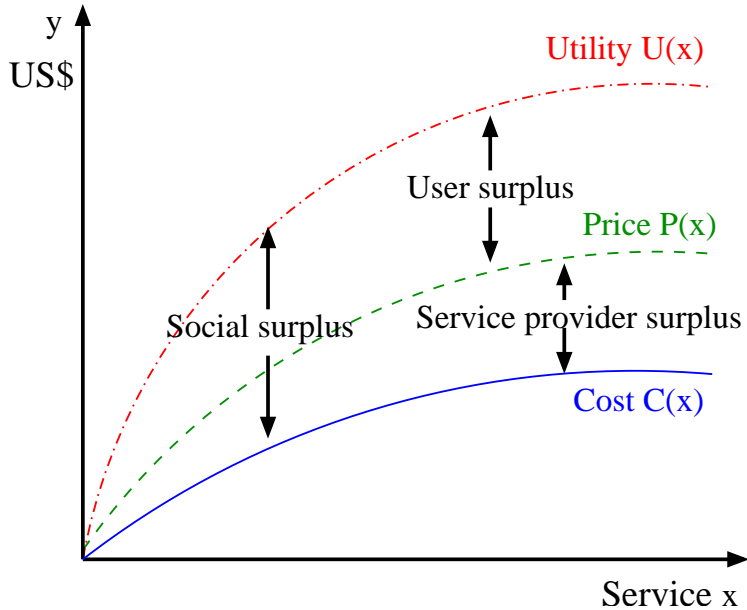


Figure 2: Utility, cost, and price functions

network infrastructure (e.g., T1, T3, or higher for virtual private networks) or are determined in some *ad-hoc* manner (e.g., the various ADSL tiers).

Usage-Sensitive Pricing. Usage-sensitive pricing policies charge users for the actual amount of traffic they generate. In current practice, ISPs charge a customer (e.g., a video-on-demand provider) based on their traffic volume using the *95-th percentile rule* [7, 21]. Specifically, the ISP measures the user’s traffic volume over 5-minute intervals during each billing period (e.g., one month), and charges the user based on the 95-th percentile value among these measured values. Typically, ISPs have a tiered pricing structure [21] in which each of the service tiers z_1, \dots, z_K , corresponds to a certain traffic volume and higher tiers are mapped to higher prices. Such a structure can be mapped to our tiered service model by considering a customer with a 95-th percentile value x such that $z_{j-1} < x \leq z_j$ as having “subscribed” to tier z_j and charging the customer accordingly¹.

3 Economic Model for Sizing of Service Tiers

In this section, we use concepts from economics to describe the relationship between users and service providers, and we propose optimization problems for selecting the service tiers. We also illustrate how to solve these problems to obtain a set of (near-) optimal tiers. In the following section, we apply Nash bargaining theory [12, 14] to determine optimal pricing strategies for this fixed (near-optimal) set of tiers.

Consider now the demand-supply relationship between the users and network service providers. On the one hand, users want to maximize the utility they obtain from the service while keeping the fee they have to pay to the service provider as low as possible; in economic terms, users want to maximize the *user surplus* [1, 4], defined as the difference between the utility they obtain from the service and the price they

¹Note that with capacity-based pricing, the tier (e.g., access speed) to which a user subscribes does not change over time (except, for instance, when a user upgrades to a higher speed), but with usage-sensitive pricing a user may be charged according to a different tier every billing period, i.e., depending on the actual traffic volume generated during each period. Nevertheless, this distinction does not affect the economic model we present in the next section.

have to pay for it. On the other hand, the network providers' objective is to charge a high fee so as to offset the cost of offering the service and make a profit; in other words, service providers want to maximize the *service provider surplus* [1,4], defined as the difference between price and cost. The concepts of user surplus and service provider surplus are illustrated in Figure 2.

From the point of view of the society as a whole, it is preferable to maximize the overall *social welfare*, defined as the sum of the user surplus plus the provider surplus (see also Figure 2). We will refer to the social welfare as *social surplus* [1,4]. Once the maximum social surplus has been determined, the users and service providers may negotiate its division into user and service provider surpluses through bargaining, as we explain in the next section.

Let us define the user surplus $S_{usr}(x) = U(x) - P(x)$, the provider surplus $S_{pr}(x) = P(x) - C(x)$, and the social surplus $S_{soc}(x) = U(x) - C(x)$. In the tiered-service network under consideration, the problems of maximizing the surplus of users, service providers, or society, amount to selecting appropriately the set of service tiers to be offered, as we discuss next.

3.1 Maximization of Expected Surplus

Let $S(x)$ be the surplus function (i.e., one of $S_{usr}(x)$, $S_{pr}(x)$, or $S_{soc}(x)$, defined above), and suppose for the moment that the vector $Z = \langle z_1, \dots, z_K = x_{max} \rangle$ of K service tiers is given. In this case (refer also to Figure 1), all users with requests in the interval² $(z_{j-1}, z_j]$ subscribe to tier z_j , incurring a surplus of $S(z_j)$, $j = 1, \dots, K$. Recalling that $f(x)$ and $F(x)$ are the pdf and cdf, respectively, of user requests, the *expected surplus* $\bar{S}(z_1, \dots, z_K)$ for the given service tier vector Z can be expressed as:

$$\begin{aligned} \bar{S}(z_1, \dots, z_K) &= \sum_{j=1}^K \left(\int_{z_{j-1}}^{z_j} S(z_j) f(x) dx \right) \\ &= \sum_{j=1}^K \left(S(z_j) \int_{z_{j-1}}^{z_j} f(x) dx \right) \\ &= \sum_{j=1}^K (S(z_j) (F(z_j) - F(z_{j-1}))). \end{aligned} \tag{1}$$

Consider now the problem of optimally selecting the service tiers *from the users' point of view*. Based on our earlier discussion, the objective of each network user is to maximize its surplus. Considering all the users in the network *as a whole*, the objective is to select the set of service tiers so as to maximize the expected aggregate user surplus, i.e., the weighted sum of the individual user surpluses in expression (1) with S_{usr} in place of $S(x)$. Similarly, the goal of the service provider is to maximize its expected aggregate surplus, while considering the welfare of the society (i.e., both users and providers), the objective would be to maximize the expected aggregate social surplus. These last two objectives are obtained by using $S_{pr}(x)$ and $S_{soc}(x)$, respectively, in place of $S(x)$ in (1).

These three optimization problems can be formally expressed as instances of the following problem which we will refer to as the *Maximization of Expected Surplus (MAX-ES)* problem. Note that the objective function (2) is nonlinear with respect to the variables z_1, \dots, z_K , hence MAX-ES is a nonlinear programming problem.

Problem 3.1 (MAX-ES) *Given the cdf $F(x)$ of user requests, an integer number K of service tiers, and a surplus function $S(x)$, find a service tier vector $Z = \langle z_1, \dots, z_K \rangle$ that maximizes the objective function*

²Note that the leftmost interval is $(z_0, z_1]$, where $z_0 = 0$ is the "null" service tier we defined earlier. Since $F(z_0) = 0$, the summation in expression (1) is correctly defined for all service tier intervals.

(expected surplus):

$$\bar{S}(z_1, \dots, z_K) = \sum_{j=1}^K (S(z_j) (F(z_j) - F(z_{j-1}))) \quad (2)$$

subject to the constraints:

$$x_{min} < z_1 < z_2 < \dots < z_K = x_{max}. \quad (3)$$

Let us assume for the moment that an optimal solution to MAX-ES can be obtained; we will discuss shortly how to find such a solution. Consider the optimal solution obtained by solving the MAX-ES problem from the perspective of users or providers (i.e., by using the user $S_{usr}(x)$ or provider $S_{pr}(x)$ surplus function in place of $S(x)$, respectively). Such a solution is unlikely to be of practical value, for two reasons. First, it assumes that users and service providers may select the service tiers optimally based only on their own interests. In reality, a service tier vector that is optimal for the users may not be acceptable to the service provider, and vice versa. Therefore, it is important to obtain a jointly optimal solution that takes into account the perspectives of both users and service providers. Second, both the user and provider surplus functions assume the existence of a pricing function $P(x)$. In general, the price function is the result of marketplace dynamics, including negotiation between users and service providers, hence it may not be known in advance.

On the other hand, the social surplus function $S_{soc}(x)$ depends only on the cost and utility functions, which are generally known in advance. Therefore, considering the welfare of the society as a whole overcomes the above difficulties since (1) it takes into account simultaneously the interests of both users (through the utility function) and providers (through the cost function), and (2) allows us to determine the optimal service tier vector without knowledge of the pricing function. Therefore, for the remainder of this paper we will consider the MAX-ES problem from the society's point of view only; although, for simplicity we will continue using $S(x)$ as the surplus function, the reader should keep in mind that from now on we assume that $S(x) = S_{soc}(x) = U(x) - C(x)$.

We also note that, from a practical point of view, the provider is the entity that determines and advertises the service tier and price structure. However, as our discussion above implies, the fact that the provider fixes the tiers does not necessarily imply that it should do this unilaterally without taking into account the users' interests (through the use of an appropriate utility function $U(x)$ and demand function $f(x)$ that can be estimated using market surveys and relevant techniques). For example, if a provider fixes the tiers and their prices at levels that do not align well with user demands and/or willingness to pay, users in a competitive environment will opt for the services of another provider.

3.2 Solution Through Nonlinear Programming

If the nonlinear objective function (2) of the MAX-ES problem is concave, and since the constraints (3) are convex, we may use the Karush-Kuhn-Tucker (KKT) conditions to find the global maximum [2]. The following lemma derives sufficient conditions for the function (2) to be concave.

Lemma 3.1 *If $S(x)$ and $F(x)$ are continuous and twice differentiable in $[x_{min}, x_{max}]$ and the two conditions*

$$S''(x)[F(x) - F(y)] + 2S'(x)F'(x) + S(x)F''(x) < 0 \quad (4)$$

$$-[S'(x)F'(y)]^2 - S(x)F''(y) \{S''(x)[F(x) - F(y)] + 2S'(x)F'(x) + S(x)F''(x)\} > 0 \quad (5)$$

are satisfied for all $x, y \in [x_{min}, x_{max}]$ with $y < x$, then the MAX-ES objective function \bar{S} is concave in the feasible area $x_{min} < z_1 < \dots < z_{K-1} < z_K = x_{max}$.

Proof. Define $\omega(x, y) = S(x)(F(x) - F(y))$. We can then rewrite (2) as:

$$\bar{S}(z_1, \dots, z_K) = \sum_{j=1}^K \omega(z_j, z_{j-1}). \quad (6)$$

Since the sum of concave functions is also a concave function, a sufficient condition for \bar{S} to be concave is for ω to be concave in the feasible area $x_{min} < y < x < x_{max}$.

The Hessian of ω is the symmetric matrix

$$\mathbf{H} = \begin{pmatrix} h_{1,1} & h_{1,2} \\ h_{2,1} & h_{2,2} \end{pmatrix}, \quad (7)$$

where

$$\begin{aligned} h_{1,1} &= \frac{\partial^2 \omega}{\partial x^2} = S''(x)[F(x) - F(y)] + 2S'(x)F'(x) + S(x)F''(x) \\ h_{2,2} &= \frac{\partial^2 \omega}{\partial y^2} = -S(x)F''(y) \\ h_{1,2} &= h_{2,1} = \frac{\partial^2 \omega}{\partial x \partial y} = -S'(x)F'(y) \end{aligned}$$

If ω is continuous and has continuous first and second derivatives, then it is concave if its Hessian is negative definite in the feasible area $x, y \in [x_{min}, x_{max}]$ with $y < x$, or:

$$h_{1,1} < 0 \quad \text{and} \quad h_{1,1}h_{2,2} - h_{1,2}^2 > 0,$$

from which the two conditions (4) and (5) follow. ■

In general, however, the objective function may not be concave. For instance, an empirically obtained cdf $F(x)$ may not be continuous, in which case the Hessian matrix is not defined. In such cases, it may be possible to formulate and solve approximate linear programming formulations, or apply branch-and-bound techniques [2]. One drawback of such solution methods is that they have to be customized to the specific cdf and surplus functions. More importantly, such methods may need a large number of iterations, or may get trapped at a local maximum.

3.3 An Efficient Approximate Solution

We now present an approximate yet efficient and accurate method for solving general instances of the MAX-ES problem. Rather than developing a sub-optimal algorithm for solving MAX-ES directly, we take a different approach: we provide an approximate formulation of MAX-ES that asymptotically converges to the formulation in (2)-(3), along with an algorithm that solves this new problem optimally.

3.3.1 An Approximate Formulation of MAX-ES

We note that it is always possible to create a discrete approximation of the pdf $f(x)$, regardless of its form, as illustrated in Figure 3. In particular, we can choose an integer $M > K$ and partition the interval $[x_{min}, x_{max}]$ into M intervals each of length equal to $\frac{x_{max} - x_{min}}{M}$. The right-hand endpoint of the m -th interval is $e_m = x_{min} + \frac{m(x_{max} - x_{min})}{M}$; we associate with e_m a discrete point mass density

$$P_m = \int_{e_{m-1}}^{e_m} f(x) dx. \quad (8)$$




Figure 3: Forming a pdf approximation: The area under $f(x)$ over an interval is paired with the right-hand endpoint of the interval.

The M pairs $\{(e_m, P_m)\}$ form the approximation of $f(x)$. We also define

$$F_m = \sum_{i=1}^m P_i, \quad m = 1, \dots, M \quad (9)$$

so that the M pairs $\{(e_m, F_m)\}$ form the approximation of the cdf $F(x)$.

In order to obtain an efficient solution to the MAX-ES problem, we also impose the additional restriction that the K service tiers may only take values from the set $\{e_m\}$ of the interval endpoints. Consequently, our objective is to solve the following discrete version of MAX-ES, which we will refer to as *Discrete-MAX-ES*.

Problem 3.2 (Discrete-MAX-ES) *Given the the M -point approximation $\{e_m, P_m\}$ of the pdf of user requests, an integer number $K < M$ of service tiers, and a surplus function $S(x)$, find a service tier vector $Z = \langle z_1, \dots, z_K \rangle$ that maximizes the objective function (approximate expected surplus):*

$$\bar{S}(z_1, \dots, z_K) = \sum_{j=1}^K (S(z_j) (F_{m_j} - F_{m_{j-1}})) \quad (10)$$

subject to the constraints:

$$z_j = e_{m_j} \in \{e_m\}, \quad j = 1, \dots, K, \quad m = 1, \dots, M \quad (11)$$

$$z_1 < z_2 < \dots < z_K = x_{max}. \quad (12)$$

Clearly, as $M \rightarrow \infty$, the pdf approximation approaches the original pdf and Discrete-MAX-ES reduces to the original MAX-ES problem.

3.3.2 Optimal Solution to Discrete-MAX-ES

Define $\Phi(m, k)$ as the optimal value of the objective function (10) when the number of intervals is m and the number of service tiers is $k \leq m$, where parameter m represents the first such intervals among the M discretized ones. Then, $\Phi(m, k)$ may be computed recursively as follows:

$$\Phi(m, 1) = S(e_m)F_m, \quad m = 1, \dots, M \quad (13)$$

$$\begin{aligned} \Phi(m, k+1) &= \max_{q=k, \dots, m-1} \{\Phi(q, k) + S(e_m)(F_m - F_q)\}, \\ &k = 1, \dots, K-1; \quad m = 2, \dots, M \end{aligned} \quad (14)$$




Figure 4: Optimal price vector

Expression (13) can be explained by observing that if there is only one tier of service, it must coincide with the right-hand endpoint of the m -th (i.e., rightmost) interval. The recursive expression (14) simply states that, for $k + 1$ service tiers, the largest tier must coincide with the right-hand endpoint of the m -th interval, and the remaining k tiers must be optimally assigned to the endpoints of any feasible interval $q, k \leq q \leq m - 1$.

The running time of the above dynamic programming algorithm to obtain $\Phi(M, K)$ is $O(KM^2)$. Note that $\Phi(M, K)$ is the value of an optimal solution to the Discrete-MAX-ES problem. The optimal solution may not be unique, in which case the algorithm will randomly return one of the optimal ones. Since Discrete-MAX-ES is an approximate formulation of the original MAX-ES problem, $\Phi(M, K)$ represents a solution close to the optimal solution to MAX-ES, regardless of the shape of the objective function (2). Clearly, the better the pdf approximation, i.e., the larger the value of M , the closer that $\Phi(M, K)$ will be to the true optimal solution for a given pdf; the tradeoff is an increase in running time. We have found that the value of $\Phi(M, K)$ converges quickly as the value of M approaches 50-100 for all the distribution functions we have considered, thus a (near-) optimal solution can be computed efficiently for any instance of MAX-ES.

After solving the MAX-ES problem, we obtain a service vector $Z^* = \langle z_1, z_2, \dots, z_K \rangle$ that maximizes the social surplus and depends only on the utility and cost functions provided by the users and network provider, respectively. Next, we describe an approach to obtaining the optimal price for each service tier in Z^* in a manner that strikes a balance between the conflicting objectives of users and providers.

4 Optimal Pricing Based on Nash Bargaining

Consider a service vector $Z^* = \langle z_1, z_2, \dots, z_K \rangle$ that maximizes the social surplus. We are interested in finding an appropriate price $P(z_j)$ for each service tier $z_j, j = 1, \dots, K$, so as to satisfy both the users and service provider. Clearly, the price for each service tier z_j should be between the values of the cost and utility functions at service level z_j , as illustrated in Figure 4.

In a free telecommunication market, the price for the service is typically the result of a negotiation process between the users and service providers. This negotiation, or bargaining, process can be thought of as a game during which each party attempts to maximize its own surplus [12–14]; the outcome of the game is




Figure 5: Relationship among the game parameters U , C , G , P_1 , P_2 , P_{high} , P_{low} , Y_1 , and Y_2

an optimal price for the service that is mutually acceptable by both parties. This game can be seen as an abstraction of market dynamics, e.g., reflecting the users' ability to compare prices from various providers and competition among providers. We also emphasize that our focus is on a bargaining game that takes place once, after which optimal prices are determined and fixed for a relatively long time compared to typical flow durations. In other words, we do not consider the scenario in which the users and/or the network attempt to set prices on a per-session basis (e.g., as in [8–10]), a scenario that we do not believe is practical. We also note that the negotiation between users and provider is implicit and captured in the game through a bargaining parameter, as explained shortly.

4.1 The Single Tier Case

Let us first consider the case $K = 1$, whereby the provider offers a single service tier z_1 . For notational convenience, in this section we will simply use U , P , and C , instead of $U(z_1)$, $P(z_1)$, and $C(z_1)$, respectively.

Let P_{high} , $P_{high} \leq U$ be the maximum price that the users would accept as a satisfactory outcome of the negotiating process (game). Similarly, let P_{low} , $P_{low} \geq C$, be the minimum price that the service provider would find acceptable. We use P_1 to denote the price that users pay for the service, and P_2 the amount that the provider receives for the service. In general, there may exist a gap G between the P_1 and P_2 ; this gap is referred to as *transaction cost* in economics. Without loss of generality, in this work we assume that the gap G has a fixed value; clearly, if $G = 0$, then $P_1 = P_2$. We also define $Y_1 = U - P_1$, and $Y_2 = P_2 - C$. Figure 5 illustrates the relationship among the game parameters U , P_{high} , P_{low} , C , P_1 , P_2 and G .

The two parties, users and service provider, are interested in dividing the *net* social surplus, i.e., the social surplus minus the transaction cost, which from Figure 5 is equal to $(U - C - G)$. As we can see, the net social surplus is the sum of Y_1 and Y_2 . Y_1 and Y_2 represent the shares of the good to be divided and stand for the excess utility (or net surplus) of users and provider, respectively. The objective is to find an optimal division of the net social surplus such that both parties feel satisfied. This optimization problem

was introduced by Nash [12, 14] as a cooperative bargaining game, and is widely used in the literature for characterizing labor negotiations and a range of other bargaining situations [17].

Let $\beta, 0 \leq \beta \leq 1$, be the bargaining power of the users, and $1 - \beta$ be the bargaining power of the service provider. Bargaining power, as defined here, refers to the relative ability of each party in the bargaining game to influence the setting of prices. Then, $\Omega = Y_1^\beta Y_2^{1-\beta}$ is the Nash product [12, 14] in the bargain. In essence, Ω is the product of the players' excess utilities, each scaled by the corresponding player's bargaining power. Our objective is to find suitable values for Y_1 and Y_2 that maximize Ω . The optimization problem can be formulated as:

$$\max_{Y_1, Y_2} \quad \Omega = Y_1^\beta Y_2^{1-\beta} \quad (15)$$

subject to the constraints:

$$Y_1 + Y_2 \leq U - C - G \quad (16)$$

$$Y_1 \geq U - P_{high} \quad (17)$$

$$Y_2 \geq P_{low} - C \quad (18)$$

The constraints (16)-(18) can be explained by referring to Figure 5.

Figure 6 plots the curve of the objective function Ω as a function of Y_1 and Y_2 . The feasible area is represented by the shaded triangle formed by the linear constraints (16)-(18). As the value of Ω increases, the curve moves upwards, and vice versa. The maximum value of Ω occurs when the curve intersects the line $Y_1 + Y_2 = U - C - G$ at exactly one point, and the coordinates of this point correspond to the optimal values for Y_1 and Y_2 . To obtain the latter values, we rewrite the optimization problem (15)-(18) in the following Lagrange form:

$$\max_{Y_1, Y_2} \quad \Omega' = Y_1^\beta Y_2^{1-\beta} + \eta(Y_1 + Y_2 - U + C + G) \quad (19)$$

where η is the Lagrange multiplier. By taking the partial derivatives of Ω' with respect to Y_1 and Y_2 , setting them equal to zero, and using the fact that $Y_1^* + Y_2^* = U - C - G$, we obtain:

$$Y_1^* = \beta(U - C - G) \quad (20)$$

$$Y_2^* = (1 - \beta)(U - C - G) \quad (21)$$

From the definition of Y_1 and Y_2 , we finally obtain the optimal prices as follows:

$$P_1^* = (1 - \beta)U + \beta(C + G) \quad (22)$$

$$P_2^* = (1 - \beta)(U - G) + \beta C \quad (23)$$

In the case of no transaction costs (i.e., $G = 0$), the price users pay is exactly the amount that the provider receives, hence:

$$P^* = P_1^* = P_2^* = (1 - \beta)U + \beta C \quad (24)$$

As we can see, the value of β determines the position of the optimal price along the line segment between the utility and cost values.

Let us now consider three special cases with respect to the value of the bargaining parameter β that provide some insight into the optimal solution of the above optimization problem.

Case 1: $\beta = 0$, i.e., the service provider has all the bargaining power; this situation arises whenever the telecommunications market is dominated by one service provider (monopoly) or a small number of providers (oligopoly). In this case we have that $P^* = U$, hence the service provider enjoys the total social surplus by squeezing out the users' surplus.




Figure 6: Optimal price point (Y_1^*, Y_2^*)

Case 2: $\beta = 0.5$, i.e., users and service provider have exactly the same bargaining power. In this case, $P^* = 0.5(U + C)$, implying that the social welfare is equally shared by the two parties.

Case 3: $\beta = 1$, i.e., the bargaining power resides solely with the users; such a scenario may arise in the telecommunications market when the supply greatly exceeds the aggregate user demand. In this case we have $P^* = C$, and the provider has to abandon any benefits (provider surplus) from supplying the service.

4.2 The Multiple Tier Case

Let us now consider the general case of $K > 1$ tiers of service. We can apply the methodology of the previous subsection to each service tier $z_j, j = 1, \dots, K$, to obtain the optimal vector of tier prices $P^* = \langle P^*(z_1), P^*(z_2), \dots, P^*(z_K) \rangle$. Let us assume for simplicity that the transaction cost G is zero; then, using expression (24) we obtain:

$$P^*(z_j) = (1 - \beta)U(z_j) + \beta C(z_j), \quad j = 1, \dots, K \quad (25)$$

Since both the utility $U(x)$ and the cost $C(x)$ are non-decreasing functions of bandwidth x , we have that:

$$P^*(z_j) < P^*(z_k), \quad 1 \leq j < k \leq K \quad (26)$$

In other words, the optimal price increases with the service tier index, i.e., with the amount of bandwidth offered to the users, consistent with intuition.




Figure 7: Cdf F_{acc} of Internet access speeds (data adapted from [3])

5 Numerical Results

To illustrate our methodology for pricing of tiered services, we consider the market for broadband Internet access under either a capacity-based or a usage-sensitive tiered pricing scheme³.

Capacity-based pricing. We have used data collected at the San Diego Network Access Point (SDNAP) and available at the CAIDA site [3] to obtain the cdf F_{acc} of Internet access speeds shown in Figure 7. We adapted the raw SDNAP data so that access speeds are in the range [256 Kb/s, 12 Mb/s], typical of current broadband speeds in the United States.

Usage-sensitive pricing. We make the assumption that the monthly amount of traffic generated by users is in the range [5MB, 1TB] and follows the bounded Pareto distribution (pdf):

$$f(x) = \frac{\alpha k^\alpha}{1 - \left(\frac{k}{p}\right)^\alpha} x^{-\alpha-1}, \quad 5 = k \leq x \leq p = 10^6, \quad 0 < \alpha < 2 \quad (27)$$

We have selected two values for parameter α , corresponding to two distribution functions:

- pdf $f_{15/85}$ has $\alpha = .00001$ and is such that approximately 15% of users generate about 85% of the total traffic, and
- pdf $f_{5/50}$ with $\alpha = .03$, for which 5% of users generate approximately 50% of the overall traffic.

The latter distribution has characteristics similar to the usage patterns reported recently by one major cable ISP [6].

For all instances of the MAX-ES problem we investigate in this study, we let the utility function be

$$U(x) = \lambda x^\gamma \log(x) \quad (28)$$

³We have conducted a large number of experiments with a range of distribution, utility, and cost functions. To avoid repetition, in this study we investigate the MAX-ES problem only with the input functions described next. Nevertheless, these input functions are characteristic of real life scenarios and the results shown are representative of what we have observed in our experiments.




Figure 8: Expected surplus against M , cdf F_{acc}

and the cost function

$$C(x) = \mu x, \quad (29)$$

hence the social surplus $S(x) = U(x) - C(x)$. The parameters of the utility function (28) can be selected such that it is an increasing, strictly concave, and continuously differentiable function of service level x ; this function has also been considered in the context of elastic traffic [19]. The parameters λ and γ can be used to control the slope of $U(x)$. In this work, we use the values $\lambda = 12$, $\gamma = 0.5$, and $\mu = 0.4$ for capacity-based pricing, and $\lambda = 9$, $\gamma = 0.5$ and $\mu = 0.05$ for usage-sensitive pricing, to ensure that the surplus function exhibits similar behavior across the different domains of the corresponding distributions.

5.1 Service Tier Selection

Let us first consider the impact of the number M of intervals in the pdf approximation (refer to Figure 3) on the convergence of the dynamic programming algorithm we presented in Section 3.3.2. Figure 8 plots the value of the optimal solution $\Phi(M, K)$ as a function of M for the cdf F_{acc} of Figure 7 and the surplus function above. Figure 9 is similar, but shows results for the Pareto cdf $F_{5/50}$.

We make two important observations from these figures. First, for a given number K of service tiers, the solution obtained by the dynamic programming algorithm $\Phi(M, K)$ converges quickly as M increases. We have run experiments with a wide range of instances of MAX-ES beyond the ones we report here, and we have found that in all cases $M = 200$ is sufficient for convergence; hence we have used this value for the experiments we present in the remainder of this section. This result confirms that the dynamic programming algorithm provides an accurate and efficient solution to the MAX-ES problem.

We also observe that the expected surplus increases with the number K of service tiers. This behavior is consistent with intuition: a larger number of tiers improves the “resolution” of the final solution and allows the dynamic programming algorithm to better tailor the tiers to the given surplus and distribution functions. The figures also demonstrate the (expected) effect of diminishing returns, as further increases in the number K of service tiers provide smaller improvements to the expected surplus.

We now compare four solutions to the MAX-ES problem in terms of the expected social surplus they achieve:




Figure 9: Expected surplus against M , Pareto cdf $F_{5/50}$

1. **Optimal:** this is the optimal dynamic programming solution to the corresponding Discrete-MAX-ES instance.
2. **Optimal-rounded:** this is the tier structure derived by rounding the values of the optimal tiers of the above solution to the nearest multiple of 256 Kb/s (for capacity-based pricing) or 10 GB (for usage-sensitive pricing). The motivation for this solution arises from considerations related to marketing the service to customers who do not have intimate knowledge of the manner in which the optimal tier structure is determined. More specifically, a tier of, say, 100 GB, is likely to seem more natural and understandable to users compared to the outcome, say, 98.54 GB, of the dynamic programming algorithm, which could well be considered arbitrary.
3. **Uniform:** the K service tiers are spread uniformly across the domain $[x_{min}, x_{max}]$, i.e.,

$$z_k = x_{min} + k \frac{x_{max} - x_{min}}{K}, \quad k = 1, \dots, K.$$

4. **Exponential:** each tier provides a level of service that is twice that of the immediately lower tier:

$$z_{k+1} = 2z_k, \quad k = 1, \dots, K - 1.$$

As a result, the tiers divide the domain $[x_{min}, x_{max}]$ into intervals of exponentially increasing length.

The uniform and exponential are simple, straightforward solutions that do not involve any optimization and are along the lines of the structures employed by major ISPs⁴. We consider them here as baseline cases and to demonstrate that they perform poorly in terms of maximizing the expected social surplus.

Since the raw expected surplus values are not comparable across different instances of the MAX-ES problem, we introduce the concept of *normalized expected surplus* to illustrate the relative performance of the four algorithms above. For a given problem instance, let S_{max} be the maximum expected surplus value achieved by any of the four algorithms over all values of the number K of tiers evaluated in our experiments.

⁴Many ADSL providers offer download speeds that follow an exponential tiering structure, e.g., 384 Kb/s, 768 Kb/s, 1.5 Mb/s, 3 Mb/s, etc. Similarly for the 5/10/20/40 GB tiers of monthly traffic used in the recent pilot program by a cable ISP [6].

If \bar{S} is the expected surplus for a given algorithm- K pair, we define the normalized expected surplus for this pair as:

$$\bar{S}_{norm} = \frac{\bar{S}}{S_{max}}, \quad 0 \leq \bar{S}_{norm} \leq 1. \quad (30)$$

The normalized expected surplus takes values in (0,1) for all instances of MAX-ES and provides insight into the relative behavior of the four algorithms.

The three figures 10-12 plot the normalized expected surplus as a function of the number K of service tiers for the three distribution functions F_{acc} , $F_{5/50}$, and $F_{15/85}$, respectively. Each figure contains four curves, each corresponding to one of the algorithms for the MAX-ES problem described above. We observe that the curves for the optimal and optimal-rounded solutions almost overlap, and exhibit the best performance by far across all the values of K except very small ones, regardless of the underlying distribution function. In particular, the exponential solution decreases rapidly for $K > 2$ to about 30-50% of the optimal expected surplus, depending on the distribution (F_{acc} or Pareto). These results demonstrate that exponential grouping of customers, though favored by ISPs, performs far from optimal from an economic standpoint. In fact, the uniform tiering structure performs better than the exponential one, but it can also be far from the optimal solution for other than very small values of K .

The main conclusion from the results shown in Figures 10-12 is that by employing our dynamic programming algorithm, which has low computational requirements, it is possible to obtain optimal tiering structures that improve the expected surplus over simple solutions by a factor of up to 2-3. More importantly, our approach makes it possible to re-optimize the tiering structures over time to accommodate evolving user demands and market conditions.

5.2 Optimal Pricing of Service Tiers

Given a vector $\langle z_1, \dots, z_K \rangle$ of K of service tiers, as well as the utility $U(x)$ and cost $C(x)$ functions, we may solve the optimization problem in Section 4 to obtain the optimal prices for each of the service tiers. Figure 13 plots the optimal prices for the $K = 5$ optimal service tiers obtained for cdf F_{acc} . As we can see, the optimal tier structure does not resemble either the uniform or exponential structures. Three price structures are shown, corresponding to the three values of the bargaining power of users $\beta = 0.25, 0.5, 0.75$. As expected, the lower the bargaining power of users, the higher the corresponding price. Also, for a fixed value of β , the price increases with the tier index, consistent with our discussion in Section 4.2. Moreover, the price increase from one tier to the next is tied directly to the shape of the utility and cost functions, thus reflecting the perspective of both users and providers.

In the previous subsection we demonstrated that the exponential and uniform tiering structures are far from optimal with respect to the expected social surplus. We now show that these structures are also suboptimal in terms of the revenue collected by the service provider. Consider a tier vector $\langle z_1, \dots, z_K \rangle$, and let $P(z_j), j = 1, \dots, K$, be the optimal price structure obtained by applying the methodology of Section 4. Then, the expected revenue \bar{R} collected by the service provider can be calculated as:

$$\bar{R}(z_1, \dots, z_K) = \sum_{j=1}^K (P(z_j) (F(z_j) - F(z_{j-1}))). \quad (31)$$

Figures 14 and 15 plot the normalized expected revenue against the number K of service tiers for the four solutions to the MAX-ES problem we described earlier; the normalized expected revenue is defined similarly to the normalized expected surplus in expression (30). Note that the highest revenue is obtained when there is only one tier, in which case all users are mapped to the highest possible service (that also incurs the highest price); such a solution is unlikely to be adopted in a market environment, and is included here for illustration purposes only. As the number K of tiers increases, the expected revenue decreases for a while and then stabilizes. The curves for the optimal and optimal-rounded solutions both converge quickly




Figure 10: Normalized social surplus comparison, cdf F_{acc}




Figure 11: Normalized social surplus comparison, Pareto cdf $F_{5/50}$




Figure 12: Normalized social surplus comparison, Pareto cdf $F_{15/85}$




Figure 13: Optimal prices, cdf F_{acc} , for $K = 5$ tiers




Figure 14: Normalized revenue comparison, cdf F_{acc}




Figure 15: Normalized revenue comparison, Pareto cdf $F_{5/50}$

to a value that is around one-half that of the maximum revenue for $K = 1$. However, the exponential and uniform solutions drop much more rapidly, eventually reaching a value that is only one-quarter (for F_{acc}) or one-sixth (for the Pareto distribution) of the maximum revenue. Again, the uniform tiering structure outperforms the exponential one, while the optimal solution achieves an expected revenue that is up to 2-3 times higher than the other two, consistent with the results of the previous section.

Since the cost of providing the service is the same regardless of what tiered structure is selected, these results indicate that by adopting simple, suboptimal solutions, the service provider may end up foregoing a substantial fraction of potential revenues. More importantly, these additional revenues are *not* at the expense of users, but rather due to the larger surplus achieved by the optimal solution. In other words, the optimal tiered structure provides substantial more value to both users and providers.

6 Extensions

We now discuss two directions in which our work may be extended.

6.1 Optimizing the Number of Service Tiers

The MAX-ES problem takes the number K of service tiers as input, and its objective function (2) is based on the assumption that there is no cost associated with offering each tier. In general, however, the total cost to the network provider of offering K tiers of service consists of two components. The first component is due to the cost of bandwidth: the higher the access speed or the amount of traffic generated by the users, the higher the cost. We used the nondecreasing function $C(x)$ to denote this cost, which may be used to represent the link cost for carrying user traffic as well as the cost of switching the traffic in the network. The second component is due to the cost of software and hardware mechanisms (e.g., queueing structures, policing mechanisms, control plane support, etc.) required inside the network for implementing a given number K of service tiers.

Let $C_t(k)$ be a nondecreasing function representing the cost of employing k service tiers. In this variant of the MAX-ES problem, the objective is to determine both the optimal number K of service tiers and their values so as to maximize the objective function:

$$\bar{S}(z_1, \dots, z_K) - C_t(K) \tag{32}$$

where $\bar{S}(z_1, \dots, z_K)$ is the expected surplus in expression (2). This problem can be solved near-optimally with the dynamic programming algorithm we presented in Section 3.3.2 after modifying the expressions (13)-(14) to account for the cost component $C_t(k)$. The running time of this algorithm is $O(M^3)$, since it has to examine all M possible values for the number of service tiers.

6.2 Tiered Structures for Bundles of Services

In this work, we have considered services that are characterized by a single parameter, e.g., access speed or amount of traffic generated. With advances in technology, more sophisticated network providers may be able to bundle and market several services in a single package. For instance, it is conceivable that providers might offer subscribers storage space on top of broadband access, or Internet connections that are characterized not only in terms of speed but also in terms of additional parameters such as availability or reliability. In such a scenario, a tiered structure would consist of service tuples, where each tuple is the set of levels of each service in the bundle corresponding to this tier. For instance, the tuple (3 Mb/s, 5 GB) would correspond to a tier in which the customer subscribes to a service bundle that offers access speed of 3 Mb/s and storage capacity equal to 5 GB. Then, the problem we considered in this work generalizes to finding a tiered service

and pricing structure that is *jointly optimal* for all the services in the bundle. This problem, which is the subject of ongoing research in our group, is further complicated by the fact that, in general, the utility and cost functions will be different for each of the bundled services

Finally, we note that the economic model we considered can be extended to include a heterogeneous population of users, in which different user segments are characterized by different utility functions, as well as heterogeneous providers (e.g., in terms of size, power of negotiation, etc). Our model only takes into account technical costs, and it is possible to extend it to consider non-technical and fixed costs as well. The implications of tiered service as the basis of a universal Internet service are also worth exploring.

7 Related Work

Pricing of Internet services using concepts from economic theory has been a subject of research for more than a decade [5, 8–11, 15, 16, 18, 20]. This is a broad area that encompasses issues from calculating the cost of resources to determining the services to offer and setting appropriate prices, and from dealing with the realities and economics of layered networks to interconnection agreements between ISPs. An initial focus was on charging as a mechanism for controlling the behavior of users, and/or for limiting usage to make room for higher paying users. Early work [9, 10] also addressed the issues of charging, rate control, and routing in communication networks carrying elastic traffic. The main finding of these studies was that the system reaches an optimum state when the network’s choice of allocated rates is at equilibrium with users’ choices of charges. The Paris Metro Pricing (PMP) scheme in [16] separates the network into independent subnetworks that behave similarly but charge their customers at different rates. A mathematical model of PMP was developed in [18] by viewing each subnetwork as a single bottleneck queue, and assuming that data packets may select the most suitable subnetwork “intelligently” by considering not only the delays but also the prices charged. The conclusion of the study was that there exist necessary and sufficient conditions for the system to attain stability. The issue of charging at the session or network layer while maintaining a clean separation between the underlying technologies was considered in [8].

More recent work has studied the issues arising in pricing multiple classes of service, especially in the context of differentiated services. A game theoretic pricing mechanism for “statistically guaranteed” service in the Internet was proposed in [20]. This mechanism was shown to offer better service and lower prices to users, and enables the provider to adopt various service and revenue models. The work in [5] also adopted a game-theoretic strategy to study a simple two-class differentiated service model, and found that the system is easy to trap into an undesirable equilibrium whenever prices do not properly reflect the quality of the service provided. Accordingly, a new dynamic pricing approach was proposed in order to avoid this problem. Finally, a free market economic model for ad-hoc wireless networks was proposed recently in [15]. Based on a greedy pricing strategy, the model maximizes the social welfare while ensuring non-negative profit for the users and service provider. This study also developed a non-greedy policy that optimizes a profit fairness metric.

Our work differs from existing literature in that our focus is on optimizing the service tiers and corresponding price structures given some information about users and providers, regardless of the underlying assumptions upon which this information is based. Consequently, our work is quite general in scope and may be applied to a variety of contexts, independently of whether the pricing scheme is capacity-based or usage-sensitive, whether charging is at the network or session/application layers, or whether the transaction is between users and provider or between providers.

8 Concluding Remarks

We proposed an economic model for tiered-service networks and developed an efficient algorithm to select the service tiers in a manner that optimizes the social surplus. We also presented a method, based on Nash bargaining, to determine the optimal price for each service tier. Our approach provides insight into the selection and pricing of Internet tiered services, as well as a theoretical framework of practical importance to network providers.

References

- [1] R. Bade and M. Parkin. *Foundations of Microeconomics*. Addison-Wesley, 2nd edition, 2004.
- [2] D. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2nd edition, 2004.
- [3] Cooperative Association for Internet Data Analysis (CAIDA). Data by bytes from SDNAP traffic. <http://www.caida.org/dynamic/analysis/workload/sdnap/>, January 2008.
- [4] B. C. Eaton, D. F. Eaton, and D. W. Allen. *Microeconomics*. Prentice Hall, 5th edition, 2002.
- [5] L. He and J. Walrand. Pricing differentiated Internet services. In *Proceedings of INFOCOM*, 195-204 2005.
- [6] C. Holahan. Time Warner’s net metering precedent. *Business Week*, June 4 2008.
- [7] C. Huang, J. Li, and K. W. Ross. Can Internet video-on-demand be profitable? In *Proceedings of ACM SIGCOMM*, pages 133–144, 2007.
- [8] N. Kausar, B. Briscoe, and J. Crowcroft. A charging model for sessions on the internet. In *European Conference on Multimedia Applications, Services and Techniques*, pages 246–261, 1999.
- [9] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
- [10] F. P. Kelly, A. Maulloo, and D. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of Operations Research Society*, 49(3):237–252, March 1998.
- [11] L. W. McKnight and J. P. Bailey (Eds.). *Internet Economics*. MIT Press, 1997.
- [12] J. F. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.
- [13] J. F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Science*, 36:48–49, January 1950.
- [14] J. F. Nash. Two-person cooperative games. *Econometrica*, 21:128–140, 1953.
- [15] M. J. Neely. Optimal pricing in a free market wireless network. In *Proceedings of IEEE INFOCOM*, 2007.
- [16] A. Odlyzko. Paris metro pricing for the Internet. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, pages 140–147, 1999.
- [17] E. Rasmusen. *Games and Information: An Introduction to Game Theory*. Blackwell Publishing, 2001.
- [18] D. Ros and B. Tuffin. A mathematical model of the Paris metro pricing scheme for charging packet networks. *Computer Networks*, 46(1):73–85, 2004.
- [19] S. Shenker. Fundamental design issues for the future internet. *IEEE Journal on Selected Areas in Communications*, 13(7):1176–1188, September 1995.
- [20] J. Shu and P. Varaiya. Pricing network services. In *Proceedings of IEEE INFOCOM*, pages 1221–1230, 2003.
- [21] H. Wang, H. Xie, L. Qiu, A. Silberschatz, and Y.R. Yang. Optimal ISP subscription for Internet multi-homing: Algorithm design and implication analysis. In *Proceedings of IEEE INFOCOM*, pages 2360–2371, 2005.