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

***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 and present an approximate yet accurate and efficient solution approach for tackling this nonlinear programming problem. 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. Our results also indicate that tiering solutions currently adopted by ISPs do not perform well.**

I. 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 or usage-sensitive. If designed and applied appropriately, multitiered pricing schemes have the potential to be a catalyst for Internet service innovation and penetration. 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. 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 services available from various providers, on the other hand, usually follow an exponential tiering structure; similar observations apply to the usage-sensitive tiered price structure employed by one cable ISP for a recent pilot program [4], which sets the tiers at 5 GB, 10 GB, 20 GB, and 40 GB of monthly download traffic.

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. We also develop 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.

Pricing of Internet services using concepts from economic theory has been a subject of research for more than a decade [3], [5]–[7], [9], [10], [12], [14]. 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. More recent work has studied the issues arising in pricing multiple classes of service, especially in the context of differentiated services. 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.

In Section II, we describe the tiered-service network, and in Section III, we introduce an economic model for selecting the set of service tiers optimally. In Section IV we use Nash bargaining theory to determine an optimal price for each of the service tiers. We present numerical results in Section V, and we conclude the paper in Section VI.

II. 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

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

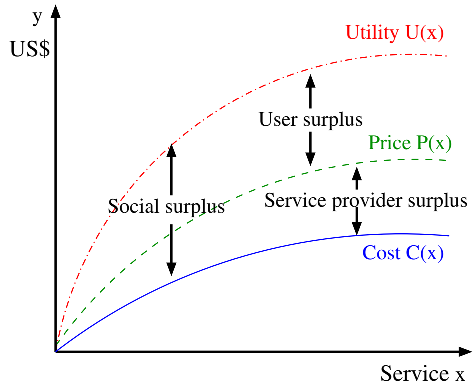

Fig. 1. Utility, cost, and price functions

distribution may be obtained empirically, or extrapolated from observed user behavior and application requirements. Let $f(x)$ and $F(x)$ be the pdf and 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. We define $Z =< z_1, z_2, \cdots, z_K >$ 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 $z_1 < z_2 < \cdots < z_K$. A user with service request $x, x_{min} \leq x \leq x_{max}$, subscribes to service tier z_j such that $z_{j-1} < x \leq z_j$. 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.

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 1. 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 1 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); however, our results are obtained for general functions $U(x)$, $P(x)$, and $C(x)$, independent of the relative behavior of the corresponding curves. We make the reasonable assumption that utility, cost, and price are all expressed in the same units (e.g., US$).

A. Applications

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.

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 $z_1, \cdots, z_K$, corresponds to a certain access speed, and users are charged based on the tier to which they have subscribed.

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* [15]. 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 [15] in which each of the service tiers $z_1, \cdots, z_K$, 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 accordingly.

III. Economic Model for Sizing of Service Tiers

Consider 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], defined as the difference between the utility they obtain from the service and the price they 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], defined as the difference between price and cost. The concepts of user surplus and service provider surplus are illustrated in Figure 1.

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 1). We will refer to the social welfare as *social surplus* [1]. 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.

A. Maximization of Expected Surplus

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

$$\bar{S}(z_1, \cdots . z_K) = \sum_{j=1}^{K} \left(S(z_j) \left(F(z_j) - F(z_{j-1}) \right) \right). \quad (1)$$

Consider now the problem of optimally selecting the service tiers *from the users' point of view*. 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., 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 $z_1, \cdots, z_K$.

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 $Z =< z_1, \cdots, z_K >$ that maximizes the objective function (expected surplus):

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

subject to the constraints:

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

Consider the optimal solution obtained by solving the MAX-ES problem from the perspective of users or providers. 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.

B. An Efficient Approximate Solution

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. In general, however, the objective function may not be concave. 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 (2)-(3), and an algorithm that solves this new problem optimally.

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. 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 $P_m = \int_{e_{m-1}}^{e_m} f(x)dx$. The M pairs $\{(e_m, P_m)\}$ form the approximation of $f(x)$. We also define $F_m = \sum_{i=1}^{m} P_i, m = 1, \cdots, M$, 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.

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 $Z =< z_1, \cdots, z_K >$ that maximizes the objective function (approximate expected surplus):

$$\bar{S}(z_1, \cdots, z_K) \quad = \quad \sum_{j=1}^{K} \left(S(z_j) \left(F_{m_j} - F_{m_{j-1}} \right) \right) \quad (4)$$

subject to the constraints:

$$z_j = e_{m_j} \in \{e_m\}, \quad j = 1, \cdots, K, \; m = 1, \cdots, M \quad (5)$$

$$z_1 < z_2 < \cdots < z_K = x_{max}. \quad (6)$$

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

2) Optimal Solution to Discrete-MAX-ES: Define $\Phi(m,k)$ as the optimal value of the objective function (4) when the number of intervals is m and the number of service tiers is $k \leq m$. Then, $\Phi(m,k)$ may be computed recursively as follows:

$$\Phi(m,1) = S(e_m)F_m, \quad m = 1, \cdots, M \tag{7}$$

$$\Phi(m,k+1) = \max_{q=k,\cdots,m-1}\{\Phi(q,k) + S(e_m)(F_m - F_q)\}, \quad k = 1, \cdots, K-1; \; m = 2, \cdots, M \tag{8}$$

Expression (7) 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 (8) 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 $O(KM^2)$. 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.

IV. Optimal Pricing Based on Nash Bargaining

Consider a service vector $Z^\star = < z_1, z_2, \ldots, z_K >$ that maximizes the social surplus. We are interested in finding an appropriate price $P(z_j)$ for each service tier $z_j, j = 1, \cdots, K$, so as to satisfy both the users and service provider. 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 [8]; the outcome of the game is an optimal price for the service that is mutually acceptable by both parties.

A. The Single Tier Case

Let us first consider the single-tier case $K = 1$. For notational convenience, 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 $Y_2 = P_2 - C$.

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 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 [8] as a cooperative bargaining game, and is widely used in the literature for characterizing labor negotiations and a range of other bargaining situations [11].

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 [8] 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 values for Y_1 and Y_2 that maximize Ω. The optimization problem can be formulated as:

$$\max_{Y_1, Y_2} \quad \Omega = Y_1^\beta Y_2^{1-\beta} \tag{9}$$

subject to the constraints:

$$Y_1 + Y_2 \leq U - C - G \tag{10}$$

$$Y_1 \geq U - P_{high} \tag{11}$$

$$Y_2 \geq P_{low} - C \tag{12}$$

Figure 2 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 (10)-(12). 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 may rewrite the optimization problem (9)-(12) in Lagrange form, from which we obtain:

$$Y_1^\star = \beta(U - C - G), \quad Y_2^\star = (1-\beta)(U - C - G) \tag{13}$$

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

$$P_1^\star = (1-\beta)U + \beta(C+G), \quad P_2^\star = (1-\beta)(U-G) + \beta C \tag{14}$$

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

$$P^\star = P_1^\star = P_2^\star = (1-\beta)U + \beta C \tag{15}$$

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 monopolized by one service provider. In this case we have that $P^\star = U$, hence the service provider enjoys the total social surplus by squeezing out the users' surplus.

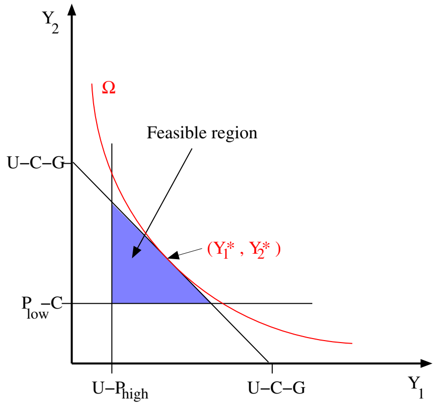

Fig. 2. Optimal price point $(Y_1^\star, Y_2^\star)$

Case 2: $\beta = 0.5$, i.e., users and service provider have exactly the same bargaining power. In this case, $P^\star = 0.5(U + C)$, and 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 when the supply greatly exceeds the aggregate user demand. In this case we have $P^\star = C$, and the provider has to abandon any benefits (provider surplus) from supplying the service.

B. 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 $z_j, j = 1, \cdots, K$, to obtain the optimal vector of tier prices $P^\star =< P^\star(z_1), P^\star(z_2), \ldots, P^\star(z_K) >$. Let us assume for simplicity that the transaction cost G is zero; then, using expression (15) we obtain:

$$P^\star(z_j) \quad = \quad (1-\beta)U(z_j) + \beta C(z_j), \quad j = 1, \cdots, K \quad (16)$$

Since both the utility $U(x)$ and the cost $C(x)$ are non-decreasing functions of bandwidth x, we have that $P^\star(z_j) < P^\star(z_k),\ 1 \le j < k \le K$. 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.

V. Numerical Results

To illustrate our methodology for pricing of tiered services, we consider both capacity-based and usage-sensitive pricing schemes.

Capacity-based pricing. We have used data collected at the San Diego Network Access Point (SDNAP) and available at the CAIDA site [2] to obtain the cdf F_{acc} of Internet access speeds. 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):

$$f(x) = \frac{\alpha k^\alpha}{1 - \left(\frac{k}{p}\right)^\alpha} x^{-\alpha-1}, \;\; 5 = k \le x \le p = 10^6, \;\; 0 < \alpha < 2 \quad (17)$$

We let $\alpha = .03$ to obtain pdf $f_{5/50}$ for which 5% of users generate approximately 50% of the overall traffic. This distribution has characteristics similar to the usage patterns reported recently by one major cable ISP [4].

For all instances of the MAX-ES problem we investigate in this study, we let the utility function be $U(x) = \lambda x^\gamma \log(x)$ and the cost function $C(x) = \mu x$, hence the social surplus $S(x) = U(x) - C(x)$. This utility function is an increasing, strictly concave, and continuously differentiable function of service level x, and has also been considered in the context of elastic traffic [13]. 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.

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

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).
3) **Uniform:** the K service tiers are spread uniformly across the domain $[x_{min}, x_{max}]$.
4) **Exponential:** each tier provides a level of service that is twice that of the immediately lower tier. As a result, the tiers divide the domain $[x_{min}, x_{max}]$ into intervals of exponentially increasing length.

The uniform and exponential are simple solutions that do not involve any optimization and are along the lines of the structures employed by major ISPs.

We use the *normalized expected surplus* metric to evaluate the performance of the four algorithms. Let S_{max} be the maximum expected surplus value achieved by any of the four algorithms over all values of K. If $\bar{S}$ is the expected surplus for a given algorithm-K pair, we define the normalized expected surplus for this pair as $\bar{S}_{norm} = \frac{\bar{S}}{S_{max}}$. This metric takes values in (0,1) and provides insight into the relative behavior of the four algorithms.

Figures 3 and 4 plot the normalized expected surplus as a function of the number K of service tiers for the distribution functions F_{acc} and $F_{5/50}$, respectively. Each figure contains four curves, each corresponding to one of the algorithms 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,

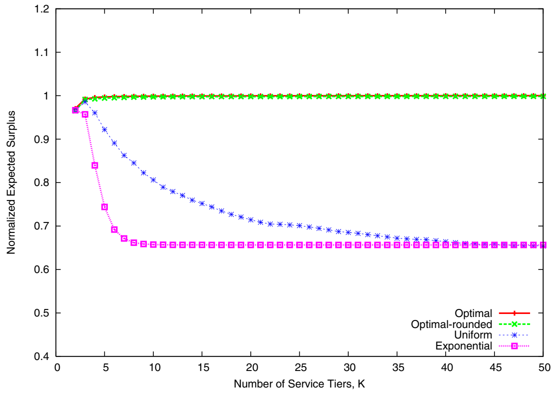

Fig. 3. Normalized social surplus comparison, cdf F_{acc}

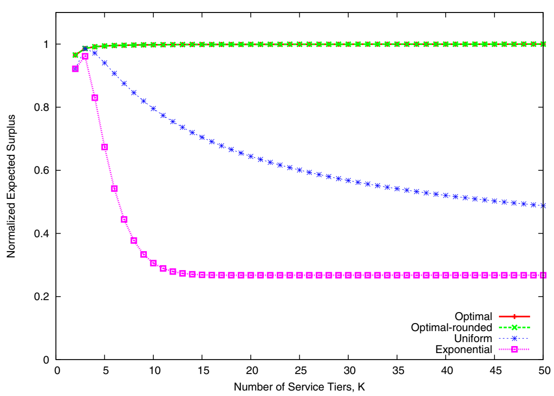

Fig. 4. Normalized social surplus comparison, Pareto cdf $F_{5/50}$

performs far from optimal from an economic standpoint. On the other hand, our dynamic programming algorithm improves over simple solutions by a factor of up to 2-3.

We now demonstrate that the exponential and uniform tiering structures are also suboptimal in terms of the revenue collected by the service provider. Consider a tier vector $< z_1, \cdots, z_K >$, and let $P(z_j), j = 1, \cdots, K$, be the optimal price structure obtained by applying the methodology of Section IV. Then, the expected revenue $\bar{R}$ collected is:

$$\bar{R}(z_1, \cdots, z_K) = \sum_{j=1}^{K} \left(P(z_j) \left(F(z_j) - F(z_{j-1}) \right) \right). \quad (18)$$

Figure 5 plots the normalized expected revenue against the number K of service tiers for the four solutions to the MAX-ES problem under cdf $F_{5/50}$. 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 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-sixth of the maximum revenue.

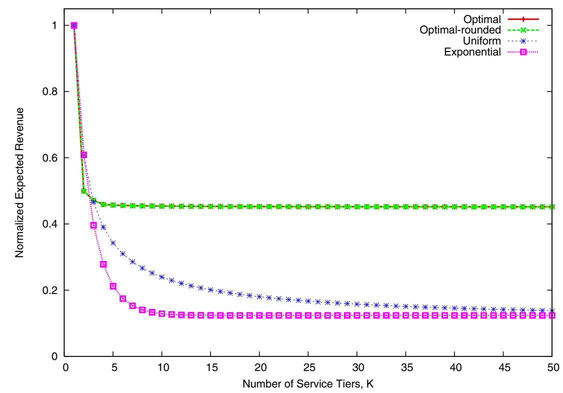

Fig. 5. Normalized revenue comparison, Pareto cdf $F_{5/50}$

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. 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.

VI. 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.

References

[1] R. Bade and M. Parkin. *Foundations of Microeconomics*. Addison-Wesley, 2nd edition, 2004.
[2] CAIDA. Data by bytes from SDNAP traffic. http://www.caida.org/dynamic/analysis/workload/sdnap/, Jan. 2008.
[3] L. He and J. Walrand. Pricing differentiated Internet services. In *Proceedings of INFOCOM*, 195-204 2005.
[4] C. Holahan. Time Warner's pricing paradox. *Business Week*, Jan. 18 2008.
[5] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
[6] 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.
[7] L. McKnight and J. Bailey (Eds.). *Internet Economics*. MIT Press, 1997.
[8] J. F. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.
[9] M. J. Neely. Optimal pricing in a free market wireless network. In *Proceedings of IEEE INFOCOM*, 2007.
[10] A. Odlyzko. Paris metro pricing for the Internet. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, pages 140–147, 1999.
[11] E. Rasumsen. *Games and Information: An Introduction to Game Theory*. Blackwell Publishing, 2001.
[12] D. Ros and B. Tuffin. A mathematical model of the Paris metro pricing scheme for charging packet networks. *Comp. Netw.* , 46(1):73-85, 2004.
[13] S. Shenker. Fundamental design issues for the future internet. *IEEE JSAC*, 13(7):1176-1188, September 1995.
[14] J. Shu and P. Varaiya. Pricing network services. In *Proceedings of IEEE INFOCOM*, pages 1221–1230, 2003.
[15] H. Wang, H. Xie, L.Qiu, A. Silberschatz, and Y.R. Yang. Optimal ISP subscription for Internet multihoming: Algorithm design and implication analysis. In *Proceedings of IEEE INFOCOM*, pages 2360–2371, 2005.