# On Optimal Tiered Structures for Network Service Bundles

Qian Lv, George N. Rouskas
Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206, USA

***Abstract*—Network operators offer a variety of tiered services in which users may select only from a small set of tiers which offer progressively higher levels of service. Service bundling, whereby several services are combined together and sold as a single package, is also common in the telecommunications market. We consider the problem of determining optimal tiering structures for service bundles using tools from economics and utility theory. Our work provides insight into the selection and pricing of Internet tiered services.**

## I. Introduction

The term *product/service bundle* refers to combining several products or services together and selling them as a single package. Product bundling is widely used as a marketing strategy [3]. Bundles are often priced at a discount to the total price that their constituent products or services would fetch if they were sold separately. Bundling can be beneficial to both consumers and sellers. The former, in addition to the lower overall price, may appreciate the lower transaction costs and simplified decision process compared with shopping for individual products or services, and may experience a better overall performance due to complementarities among the bundle components. For sellers, bundling has the potential to reduce production and transaction costs, reduce customer churn, and increase revenue and profitability. In particular, bundling is most successful as a marketing strategy whenever the marginal costs of bundling are low, customer acquisition costs are high, and there are economies of scale in production and distribution of the bundled products. Consequently, bundling is common in industries that share these characteristics, including the telecommunications and cable TV industry, the software business, and the fast food industry, among others.

Network operators have also developed a variety of *tiered service* models in which users may select only from a small set of service *tiers* which offer progressively higher levels of service. Service bundles, and associated tiered structures, are prevalent in the telecommunications market. For instance, wireless providers combine voice, data, and text services into tiered subscription packages marketed to users, where a given tier corresponds to a certain combination of values for voice minutes, Internet data, and text messages available to the user during the billing period. Similarly, ISPs may bundle a broadband access service with an email or web hosting service (for which fees may be based on the amount of traffic handled), and possibly an online storage service (characterized by the amount of data the user may store on the provider's servers).

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

In [6] we investigated the benefits of tiered service in the context of MPLS networks. In particular, we considered the problem of "sampling" the continuous range of possible rates to select a small set of discrete bandwidth levels (tiers) that are made available to users, and we presented a sophisticated dynamic programming algorithm of linear complexity to obtain the optimal set of bandwidth tiers. The main contribution of this study was to demonstrate that the benefit of offering a small, predetermined set of tiers rather than supporting arbitrary rates over a large continuous range comes almost for free, as the performance degradation (e.g., in terms of call blocking) compared to a continuous-rate network is negligible. In more recent work [5] we developed an economic model for reasoning about and pricing Internet tiered services.

In [5], [6] we considered a single network service. With the proliferation of network service bundles, it is desirable to design multi-dimensional tiered structures, where each dimension corresponds to a certain level of one distinct service in the bundle. Note that tiering is even more important for service bundles since the space of potential service levels grows as the product of the space for each service component. The objective in this case would be to determine a tiered structure that is *jointly optimal* for a vector of network services. With such a tiered structure, a user with a certain level of requirements for each service component would subscribe to the tier that offers a level at least equal to its requirements across all dimensions of service. In [4] we modeled this problem as a directional $p$-median problem in multiple dimensions and we showed it to be NP-complete. We also employed concepts from location theory to develop efficient algorithms that construct near-optimal tiering structures for service bundles given some information regarding the user demands and the cost to the provider for providing the services.

In this paper we consider the problem of determining optimal tiering structures for service bundles using tools from economics and utility theory. The paper is organized as follows. In Section II we develop an economic model for bundled network services, we introduce the Cobb-Douglas utility, and formally define the problem of selecting jointly the tiers and their prices so as to maximize the expected profit (i.e., provider surplus [2]) of the ISP under user budget constraints. In Section III we develop dynamic programming algorithms both for the case of predetermined tiers (i.e., when only price is subject to optimization) and the general version of the problem. We present numerical results in Section IV, and we conclude the paper in Section V.

## II. Economic Model of Service Bundling

Consider an ISP that offers two services. One service, characterized by parameter $x$ (e.g., access speed), may be offered at levels between a minimum $x_{min}$ and a maximum $x_{max}$. The second service, say, web hosting, is also characterized by a single parameter $y$ (e.g., corresponding to monthly amount of traffic handled), with $y$ also taking values between a minimum $y_{min}$ and a maximum $y_{max}$ level. The ISP bundles the two services into a package, and offers a tiered structure with $p$ tiers for the combined service. We let $Z = \{(z_1, t_1), \ldots, (z_p, t_p)\}$ denote the set of $p$ distinct service tiers, where the $j$-th tier $(z_j, t_j), t = 1, \ldots, p$, corresponds to an amount $z_j$ for service $x$ and an amount $t_j$ for service $y$.

We let $C(x, y)$ denote the cost to the ISP of offering a service bundle $(x, y)$ of the two services. We also let $P(z_j, t_j), j = 1, \ldots, p$, denote the price that the ISP charges subscribers to tier $(z_j, t_j)$. Without loss of generality, we assume that tiers are labeled such that

$$P(z_{j-1}, t_{j-1}) \quad < \quad P(z_j, t_j), \quad j = 2, \ldots, p. \tag{1}$$

For mathematical convenience, we also define the "null" service tier $(z_0 = 0, t_0 = 0)$ with price $P(z_0, t_0) = 0$, as well as a fictitious $(p+1)$-th service tier such that $P(z_{p+1}, t_{p+1}) = \infty$.

The value that users receive from a bundle $(x, y)$ of the two services is described by the utility function $U(x, y)$. In essence, the utility function imposes a pairwise ranking of bundles by order of preference, where *preference* is a transitive relation. More precisely, if $U(x, y) > U(x', y')$, then bundle $(x, y)$ is said to be strictly preferred to bundle $(x', y')$. On the other hand, if $U(x, y) = U(x', y')$, the two bundles are equally preferred, and the consumer is said to be *indifferent* between the two bundles. In particular, a curve

$$U(x, y) \quad = \quad u \tag{2}$$

is referred to as an *indifference curve* since the user has no preference for one bundle over another among the bundles represented by points along this curve. In other words, each point on an indifference curve provides the same level of utility (value, or satisfaction) to the user. Indifference curves are typically used to represent demand patterns for product or service bundles observed over a population of consumers.

Fig. 1 shows a set of indifference curves, each associated with a different utility level. In this figure, utility is measured along the $z$ (vertical) axis, and the indifference curves are simply the projections of the function $U(x, y) = u$, for various values of constant $u$, on the $xy$ plane. In Fig. 1, users would rather be on curve $I_7$ rather than $I_6$; they would also rather be on curve $I_6$ rather than on $I_5$, and so on, but they do not care where they are on a given indifference curve. Indifference curves are similar to topographical maps, in that each point along a given curve is at the same "altitude" above the floor.

The characteristics of the curves in Fig. 1 are typical of indifference curves in general. Specifically, indifference curves are defined only on the positive quadrant of the $xy$ plane, and they are negatively sloped and convex; in other words,

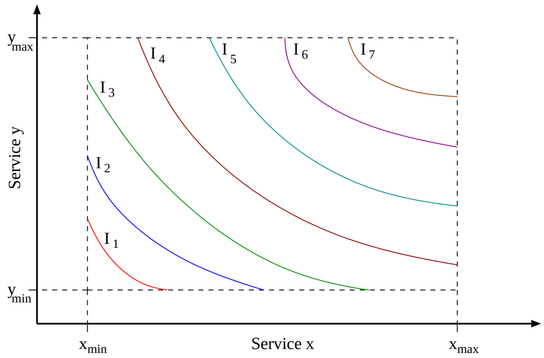


Fig. 1. Indifference curves, $I_1, \ldots, I_7$, such that $U(x, y) =$ constant along each curve (utility is measured on the vertical $z$ axis)

as the quantity of one service or good $x$ (respectively, $y$) that is consumed increases, it must be offset by a decrease in the quantity consumed of the other good $y$ (respectively, $x$), so as to keep utility (satisfaction) constant.

The Cobb-Douglas family of functions [1] generate indifference curves with the characteristics shown in Fig. 1 and are widely used as utility functions in this context. This parameterized family of functions is defined as:

$$U(x, y) \quad = \quad x^{\alpha} y^{1-\alpha}, \quad 0 \leq \alpha \leq 1, \tag{3}$$

where $\alpha$ is a parameter whose value is used to specify a certain function within the family. Then, the indifference curve for a constant level $u$ of utility is given by:

$$y \quad = \quad u^{\frac{1}{1-\alpha}} x^{\frac{-\alpha}{1-\alpha}} \tag{4}$$

We will use the Cobb-Douglas utility in (3) as the utility function in this paper.

We make the assumption that each user has a budget $B$, where $B$ is a random variable defined in the interval $[B_{min}, B_{max}]$. We let $f(B)$ and $F(B)$ denote the PDF and CDF, respectively, of random variable $B$. We make the assumption that a consumer will make a purchase if and only if the price of the product is no greater than the consumer's budget. More specifically, given a set $Z$ of $p$ tiers and a price structure consistent with (1), a user will subscribe to the tier $(z_j, t_j)$ with the highest index $j$ whose price $P(z_j, t_j)$ does not exceed the user's budget $B$.

We are interested in selecting a set of service tiers for the bundled services, and determining their prices, so as to maximize the expected provider surplus (i.e., profit). We call this the *maximization of expected provider surplus in two dimensions (MAX-ES-2D)* problem, defined formally as:

*Problem 2.1 (MAX-ES-2D):* Given the cost and utility functions $C(x, y)$ and $U(x, y)$, respectively, defined in the domain $[x_{min}, x_{max}] \times [y_{min}, y_{max}]$, and the CDF $F(B)$ of user budgets, find a set $Z = \{(z_1, t_1), \ldots, (z_p, t_p)\}$ of $p$ service tiers and their respective prices $P(z_j, t_j)$ that maximizes the following objective function representing the expected provider surplus:

$$\bar{Q}_{pr}(Z) = \sum_{j=1}^{p}((P(z_j,t_j) - C(z_j,t_j)) \times (F(P(z_{j+1},t_{j+1})) - F(P(z_j,t_j)))) \quad (5)$$

under the constraints:

$$P(z_1,t_1) < P(z_2,t_2) < \dots < P(z_p,t_p) \quad (6)$$

$$P(z_j,t_j) \le U(z_j,t_j), \quad j = 1,\dots,p \quad (7)$$

$$x_{min} \le z_j \le x_{max}, \; y_{min} \le t_j \le y_{max}, \; j = 1,\dots,p \quad (8)$$

Note that the terms $F(P(z_{j+1},t_{j+1})) - F(P(z_j,t_j)), j = 1,\dots,p$, in the right-hand side of (5) represent the fraction of users whose budgets fall in the intervals $[P(z_j,t_j), P(z_{j+1},t_{j+1}))$, hence they will subscribe to tier $j$ (recall also that we have defined $P(z_{p+1},t_{p+1}) = \infty$, and that $F(P(z_{p+1},t_{p+1})) = 1$). Also, constraint (7) states that the price of a service tier has to be no greater than the utility (value) of this tier to users, since otherwise users will not subscribe even if their budget allows them to do so.

We have the following result.

*Lemma 2.1:* Let $Z = \{(z_1,t_1),\dots,(z_p,t_p)\}$ be an optimal solution to MAX-ES-2D. Let $u_j = U(z_j,t_j), j = 1,\dots,p$. Then, for all $j$, tier $(z_j,t_j)$ is the point on the indifference curve $U(x,y) = u_j$ that minimizes the cost $C(x,y)$.

**Proof.** By contradiction. Assume that in the optimal solution the $j$-th tier is such that $C(z_j,t_j)$ is not the minimum cost point on the indifference curve $U(x,y) = u_j$. Let $(z'_j,t'_j)$ be such a minimum cost point, and let $Z'$ be the solution derived from $Z$ with $(z_j,t_j)$ replaced by $(z'_j,t'_j)$. Since the utility and price of the $j$-th tier is not affected by this change, from (5) it is clear that $\bar{Q}_{pr}(Z') > \bar{Q}_{pr}(Z)$, contradicting the assumption that $Z$ is an optimal solution. ■

## III. Approximate Solution to MAX-ES-2D

### *A. The Fixed Tier Case*

Consider first a special variant of the MAX-ES-2D problem in which the $p$ service tiers are predetermined and part of the input, and not subject to optimization; this variant arises in the case of the uniform and exponential tiering structures that we introduce in Section IV. The cost of each tier is completely determined in this case, and for simplicity we let $C_j = C(z_j,t_j), j = 1,\dots,p$. The price of each tier $j$ is equal to the utility, i.e., $P_j = P(z_j,t_j) = U(z_j,t_j)$, and we let $P_0 = P(x_{min},y_{min}) = U(x_{min},y_{min})$. Hence, the provider's surplus for a fixed tier structure $Z$ can be obtained from the following expression:

$$\bar{Q}(Z) = \sum_{j=0}^{p-1}(P_j - C_j)(F(P_{j+1}) - F(P_j)) \quad (9)$$

### *B. Cost Minimization on an Indifference Curve*

Before we tackle the general version of the MAX-ES-2D problem, we note that, because of Lemma 2.1, each tier in an optimal solution is the point on an indifference curve with the minimum cost among all points on this curve. Therefore, let us consider the optimization problem of the form:

$$\text{Minimize } \; C(x,y) \; \text{ subject to } \; U(x,y) = u. \quad (10)$$

Depending on the form of the cost and utility functions, this problem may be solved exactly or approximately using standard optimization techniques. Here we will only consider cost functions $C(x,y)$ that are linear functions of $x$ and $y$:

$$C(x,y) = c_1 x + c_2 y. \quad (11)$$

Assuming Cobb-Douglas utility functions as in (3), we may solve for $y$ as a function of $x$:

$$y = \left(\frac{u}{x^\alpha}\right)^{1/(1-\alpha)} \quad (12)$$

Substituting this value of $y$ into the cost function (11), we obtain an expression for the cost that is a function of $x$ only:

$$C(x) = c_1 x + c_2 \left(\frac{u}{x^\alpha}\right)^{1/(1-\alpha)}. \quad (13)$$

The first and second derivatives of $C(x)$ are:

$$C'(x) = c_1 + c_2 \, u^{\frac{1}{1-\alpha}} \left(\frac{-\alpha}{1-\alpha}\right) x^{\frac{-1}{1-\alpha}} \quad (14)$$

$$C''(x) = c_2 \, u^{\frac{1}{1-\alpha}} \frac{\alpha}{(1-\alpha)^2} x^{\frac{\alpha-2}{1-\alpha}}. \quad (15)$$

If there are no other constraints, we can just let $C'(x) = 0$, and obtain the optimal values:

$$x^\star = u\left(\frac{c_1(1-\alpha)}{c_2\alpha}\right)^{\alpha-1}, \quad y^\star = u\left(\frac{c_1(1-\alpha)}{c_2\alpha}\right)^{\alpha}. \quad (16)$$

It is easy to prove that $C''(x^\star) > 0$. Thus, $C(x)$ achieves its minimum value at $x^\star$, hence the original cost function $C(x,y)$ is minimized at $(x^\star,y^\star)$.

Recall, however, that $x$ and $y$ are defined only between respective minimum and maximum values. Consider the above optimization problem under the additional constraints:

$$x_{min} \le x \le x_{max} \qquad y_{min} \le y \le y_{max}. \quad (17)$$

It easy to see that when $x < x^\star$, $C'(x) < 0$, and when $x > x^\star$, $C'(x) > 0$. Consequently, whenever the unconstrained minimum point $(x^\star,y^\star)$ from (16) lies outside the feasible region defined by constraints (17), the minimum point within the feasible region can be obtained as follows:

- if $x^\star > x_{max}$, then $x^\star = x_{max}$, and $y^\star = \left(\frac{u}{x^\alpha_{max}}\right)^{\frac{1}{1-\alpha}}$;
- if $x^\star < x_{min}$, then $x^\star = x_{min}$, and $y^\star = \left(\frac{u}{x^\alpha_{min}}\right)^{\frac{1}{1-\alpha}}$;
- if $y^\star > y_{max}$, then $x^\star = \left(\frac{u}{y^{1-\alpha}_{max}}\right)^{\frac{1}{\alpha}}$, and $y^\star = y_{max}$; and
- if $y^\star < y_{min}$, then $x^\star = \left(\frac{u}{y^{1-\alpha}_{min}}\right)^{\frac{1}{\alpha}}$, and $y^\star = y_{min}$.

### C. Service Tier Optimization

The most general version of the MAX-ES-2D problem involves the selection of service tiers and their respective prices so as to maximize provider surplus, subject to the constraints (6)-(8). The utility function $U(x, y)$ provides a relative ranking of service bundles $(x, y)$ in terms of user preference, and the utility of any service tier will lie in the interval $[U_{min}, U_{max}]$, where $U_{min} = U(x_{min}, y_{min})$ and $U_{max} = U(x_{max}, y_{max})$. Therefore, the problem can be logically decomposed into two subproblems:

1) find the indifference curve $I_j$ (i.e., utility value $u_j \in [U_{min}, U_{max}]$) on which each optimal service tier $(z_j, t_j)$ lies and set the price of the tier to $u_j$; and
2) set tier $(z_j, t_j)$ to the point in indifference curve $I_j, j = 1, \ldots, p$, that minimizes the provider cost $C(z_j, t_j)$.

The second subproblem was addressed in the previous subsection. Next, we develop a dynamic programming solution for the first subproblem.

To this end, we employ a discretization technique. Specifically, we divide the domain $[U_{min}, U_{max}]$ of the utility function $U(x)$ into $K > p$ equal-length sub-intervals, such that the right endpoint $U_k$ of the $k$-th sub-interval is $U_k = \frac{k(U_{max} - U_{min}}{K}, k = 1, \ldots, K$. We also restrict the tiers to take values only from the discrete set $\{U_k\}$ of indifference curves, rather than the continuous set $[U_{min}, U_{max}]$. Let $\Upsilon(k, l, w)$ denote the optimal value of (5) when there are $k$ sub-intervals, $l$ tiers and the $l$-th tier is set at the indifference curve of utility value $U_w, w \le k$. Let also $C_w^\star, w = 1, \ldots, K$, denote the minimum cost on the indifference curve of utility $U_w$. Then, we may write the following recursion:

$$\Upsilon(k, 1, w) = (U_w - C_w^\star)(F(U_k) - F(U_w)), \quad k = 1, \ldots, K; w = 1, \ldots, k \qquad (18)$$

$$\Upsilon(k, l+1, w) = \max_{q=l,\ldots,w} \Big\{(U_w - C_w^\star)(F(U_k) - F(U_w)) + \max_{v=l,\ldots,q} \{\Upsilon(q, l, v)\}\Big\}$$
$$l = 1, \ldots, p-1; \quad k = 2, \ldots, K; w = 1, \ldots, k. \qquad (19)$$

Expression (18) can be explained by noting that when there are $k$ sub-intervals and only one tier with a price set to $U_w$, the customers who subscribe to the service at this price are those with budgets equal to or greater than $U_w$, or a fraction $(F(U_k) - F(U_w))$ of the total user population. For each subscriber, the provider has a profit of $U_w - C_w^\star$, hence the expected surplus is given by (18). Expression (19) can be similarly explained. Once $\Upsilon(k, l, w)$ has been computed for all values of $k$, $l$, and $w$, the overall optimal for $p$ tiers and $K$ intervals can be determined as:

$$\max_w \ \Upsilon(K, p, w) \qquad (20)$$

The overall running time complexity of this dynamic programming algorithm is $O(pK^4)$.

As $K \to \infty$, this discrete version of MAX-ES-2D approaches the original version in which the tiers are continuous variables. We have conducted a large number of experiments (omitted due to space constraints) which indicate that $K = 100$ is sufficient for the dynamic programming algorithm to converge; hence, we use this value in the performance study we present in the next section.

## IV. Numerical Results

In order to evaluate tiering structures for service bundles, we consider an ISP offering a bundle of two services, namely, access speed $x$ and web hosting traffic handled $y$. The domain of service $x$ is [256 Kbps, 12 Mbps], while the domain of service $y$ is [100 MB, 1 TB]. We consider the following tiering structures in our study:

1) **Optimal:** the set of tiers $Z = \{(z_1, t_1), \ldots, (z_p, t_p)\}$ obtained as a solution to the dynamic programming algorithm (18)-(20), where $z_i \in$ [256 Kbps, 12 Mbps] and $t_i \in$ [100 MB, 1 TB].
2) **Optimal-rounded:** the set of tiers obtained after rounding the values of each tier $(z_i, t_i) \in Z$ such that $z_i$ is rounded to the nearest multiple of 256 Kbps and $t_i$ is rounded to the nearest multiple of 100 MB.
3) **Uniform-uniform:** the tier structure constructed by (1) obtaining a uniform tiering structure $\{z_1, \ldots, z_p\}$ for service $x$ by spreading the $p$ tiers across the domain [256 Kbps, 12 Mbps], (2) obtaining a uniform structure $\{t_1, \ldots, t_p\}$ for service $y$ by spreading the $p$ tiers across the domain [100 MB, 1 TB], and (3) pairing the tiers of same index in the two sets to form the tiers $\{(z_1, t_1), \ldots, (z_p, t_p)\}$ for the bundle.
4) **Exponential-exponential:** this tier structure is obtained in a similar manner as uniform-uniform, except that the $p$ single-service tiers divide their respective domain into exponential intervals (i.e., intervals that double in length, from left to right).
5) **Uniform-exponential:** the tier structure in which $p$ uniform (respectively, exponential) tiers are obtained for service $x$ (respectively, service $y$), which are then paired to obtain the $p$ tiers for the service bundle.
6) **Exponential-uniform:** the tiers for service $x$ are exponential and those of service $y$ are uniform.

Note that uniform and exponential tiered structures are similar to those employed by major ISPs (e.g., ADSL tiers of 768 Kbps, 1.5 Mbps, 3 Mbps, 6 Mbps, etc). For the last four tiering solutions, the $p > 1$ service tiers are fixed. Therefore, the provider surplus in this case was obtained from expression (9).

We use the Cobb-Douglas utility function in expression (3) with parameter $\alpha = 0.6$, and a linear cost function as in expression (11), with $c_1 = 0.1$ and $c_2 = 0.01$; these values for $c_1$ and $c_2$ were selected so that neither term of the cost function dominates across the domains of services $x$ and $y$. Plots of the utility and cost functions are shown in Fig. 2.

In order to study the effect of the distribution of user budgets, we consider three distinct distributions in the domain $[B_{min} = 10, B_{max} = 1000]$:

- a *decreasing* distribution, $f(B) = -\frac{2B}{(B_{max} - B_{min})^2} + \frac{2B_{max}}{(B_{max} - B_{min})^2}$, with mean 345, in which the mass of the

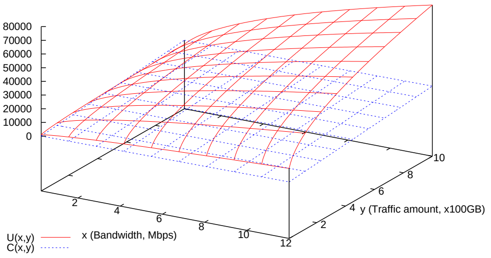


Fig. 2. Utility and cost functions used for the experiments

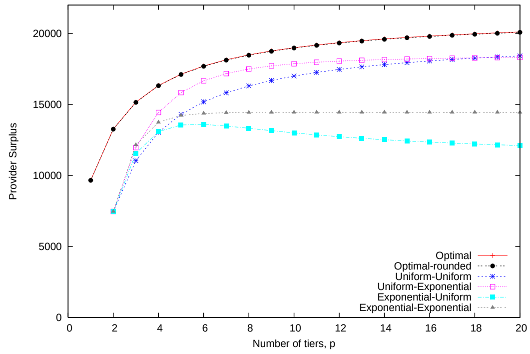


Fig. 3. Tiered structure comparison, decreasing budget distribution

distribution is concentrated at lower budget values (less affluent population),

- a *uniform* distribution with PDF $f(B) = \frac{1}{B_{max}-B_{min}} = \frac{1}{990}$ and mean 495, and
- an *increasing* distribution, $f(B) = \frac{2B}{(B_{max}-B_{min})^2} - \frac{2B_{min}}{(B_{max}-B_{min})^2}$, with mean 650, in which the mass of the distribution is concentrated at higher budget values (more affluent population).

Figs. 3-5 plot the expected provider surplus for the decreasing, uniform, and increasing, respectively, distribution of user budgets. Each figure shows six curves, corresponding to the six tiered structures above. A first observation is that, for a given tiered structure and a given number of tiers, the expected provider surplus depends directly on the distribution of user budgets. Specifically, the provider surplus increases from Fig. 3 (decreasing distribution) to Fig. 4 (uniform distribution) to Fig. 5 (increasing distribution). This result is directly due to the fact that the average user budget is lowest under the decreasing distribution and highest under the increasing distribution.

We also observe that the optimal and optimal-rounded structures outperform the other four fixed-tier structures. The optimal-rounded curves lie a little lower than the corresponding optimal curves, as a result of rounding in two dimensions. More importantly, structures which include exponential tiering of at least one service are the worst performers in terms of provider profits. This behavior demonstrates that exponential tiers currently favored by major ISPs are far from optimal. Overall, these results provide a strong indication that the optimization methodology we developed in this paper represents a valuable tool for service providers.

## V. Concluding Remarks

We have investigated tiered structures for bundles of network services with the objective of maximizing provider profits under user constraints. We have developed an efficient dynamic programming algorithm for determining jointly the service tiers and their prices. Although we only considered bundles of two services, our work may be extended to bundles of more than two services.

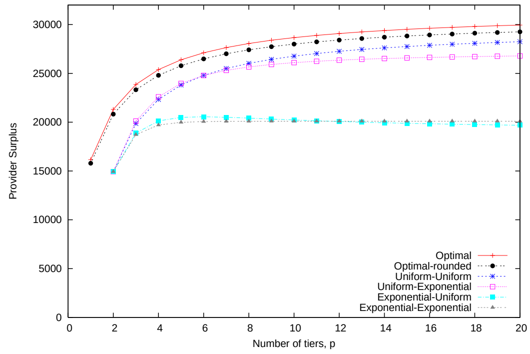


Fig. 4. Tiered structure comparison, uniform budget distribution

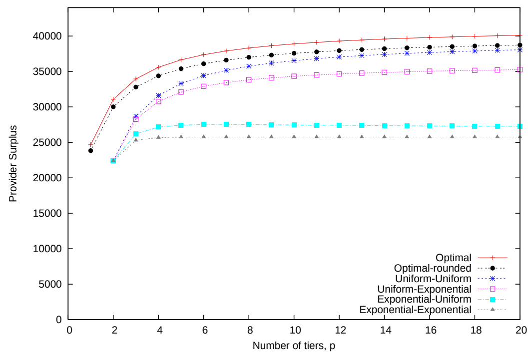


Fig. 5. Tiered structure comparison, increasing budget distribution

## References


[1] H. L. Ahuja. *Modern Microeconomics: Theory and Applications*. S. Chand and Company Ltd., 2004.

[2] R. Bade and M. Parkin. *Foundations of Microeconomics*. Addison-Wesley, 2nd edition, 2004.

[3] G. R. Bitran and J. C. Ferrer. On pricing and composition of bundles. *Production and Operations Management*, 16(1):93–108, February 2007.

[4] L. E. Jackson, G. N. Rouskas, and M. F. M. Stallmann. The directional $p$-median problem: Definition, complexity, and algorithms. *European Journal of Operations Research*, 179(3):1097–1108, June 2007.

[5] Q. Lv and G. N. Rouskas. An economic model for pricing tiered-service networks. In *Proceedings of IEEE ICC 2009*, June 2009.

[6] G. N. Rouskas and N. Baradwaj. A framework for tiered service in MPLS networks. In *Proceedings of IEEE INFOCOM 2007*, pages 1577–1585, May 2007.