

Pricing Resources on Demand

Costas Courcoubetis*, Sergios Soursos* and Richard Weber†

*Athens University of Economics and Business

Email: {courcou, sns}@aueb.gr

†University of Cambridge

Email: rrw1@statslab.cam.ac.uk

Abstract—Traditional contracts for network and computing resources are of “static” type where the customer is buying the right to use for a given price a fixed amount of resources for a long period of time. Typical examples are the case of contracting bandwidth in access networks and VPNs and the case of computing infrastructure that a customer leases (or buys) for fulfilling its IT needs. Current technology in access networks and Grid computing allows suppliers to offer more flexible contracts to their customers allowing them to choose dynamically the amount of resources they are allowed to use at a given time. This flexibility may benefit the customers with bursty demand since it allows them to obtain resources only when they need them and pay only when they use them. We define contracts where time is discrete and a customer is allowed to buy a fixed amount of resources ahead of time for a price a , the “static” part of the contract, and complement this at each new time period by purchasing an extra amount at price b , the “dynamic” part of the contract. We investigate the properties of such contracts and compare them with contracts of purely static or dynamic type. Our results suggest that in general suppliers and customers are both better off when using such mixed contracts, and that purely dynamic contracts may not always be preferable compared to purely static ones. We also show that under price competition of suppliers using static contracts against suppliers using dynamic contracts, at the equilibrium both suppliers may secure some profit by segmenting the market.

I. INTRODUCTION

Traditional bandwidth contracts have been for fixed bandwidth pipes and are long term, typically a year. We are thinking of a network supplier that leases high capacity broadband lines to other large companies, such as banks or Internet service suppliers (ISPs). A bank might lease a 155Mbps line between two locations for one year. An AOL subscriber might sign a year’s contract for the service known as ‘AOL Broadband Gold’ (2Mbps). Typically, such Internet access or VPN services do not offer any flexibility to the customer to alter dynamically the size of the pipe. Similarly, when acquiring computing resources such as servers and PCs, a company’s IT department usually makes yearly leasing contracts for fixed amount of computing resources, or buys these resources which is the equivalent of a fixed size contract for an even longer period of time. The problem customers face in making such contracts is that their demand for resources is not constant but bursty and in many cases unpredictable. They hence face the

This work is part of the project “Economic and Architectural Aspects for Services of Differentiated Quality in high-speed Broadband and Wireless Access Networks”, co-funded by the “European Social Fund & National Resources - EPEAEK II - IRAKLITOS Fellowships for research of Athens University of Economics and Business”. This work is also supported by the IST project GridEcon (IST-2006-033634) funded by the EU.

risk of acquiring for a long time period resources that may not be effectively used, or acquiring fewer resources that will fail to meet high short term demand and peak load requirements.

Emerging technology in access networks and Grid computing allows suppliers to offer more flexible contracts to their customers providing them the ability to choose dynamically the amount of resources they are allowed to use at a given time. For instance, DSLForum’s BroadbandSuite platform [1] and IPsphere’s framework [2] propose solutions that will offer the ability to alter the size of the bandwidth pipe with which a customer is provided on a time scale that is much finer than yearly. They also provide accounting and billing for such dynamic services. Grid computing offers a direct analogy to bandwidth provisioning. Current Grid utility computing architectures such as SUN Grid Compute Utility [3] allow customers to buy computing cycles on demand while also using their existing computing facilities. This flexibility should benefit the customers with bursty demand since it allows them to obtain resources only when they need them and pay only when they use them. If priced correctly, these services may increase the revenue of the provider by obtaining a share of the added value to his customers.

In this paper, we investigate the properties of a new type of contract, the “mixed” contract, which is composed by combining static and dynamic contracts. Under a mixed contract a customer is allowed to buy a fixed amount of resources ahead of time for a price a , the “static” part of the contract, and complement this at each new time period by purchasing an extra amount at price b , the “dynamic” part of the contract.

Dynamic contracts or the dynamic part of a mixed contract should not be confused with *dynamic pricing approaches* which aim to control network traffic by means of prices. One such approach is *peak-load pricing* [4], [5] where prices can adjust to the fluctuations of demand and may reflect the investment made in order to serve the high-demand periods. Due to the fact that mixed contracts have two parts, they may be considered as *two-part tariffs*. In the literature, two-part tariffs are used by the providers to obtain a greater portion -if not all- of a customer’s utility. The static part of such a tariff is a lump sum that the user must pay in order to have the right to use the resources [6] or a way for the provider to cover his fixed costs. In some cases, this lump sum may give the right to use a small amount of resources. But in our pricing scheme, no lump sum and no fixed amount of resources is introduced. Our two-part contract aims to make more flexible the customer in the way he expresses his bandwidth needs. Throughout the paper we consider that both prices (a and b) are fixed, i.e. they

are not affected by the demand. On the other hand, these prices can be considered as usage-based prices since customers are charged based on the volume of bandwidth they consume.

An interesting question we analyze in this paper concerns the position of the suppliers of dynamic services. Will they be always better off by providing such services or they may cannibalize their lucrative static contracts? Will bandwidth consumption increase or decrease? What will be the result of price competition between dynamic and static service suppliers? Will one supplier get the whole market share by displacing the other? We show that many interesting facts may occur. For instance, in a monopolist situation, a customer may not always be better off when the supplier offers purely dynamic contracts. Also our models suggest that resource consumption may drop when dynamic contracts are used (which justifies the fears of the network operators for allowing dynamic contracts). Moreover, we show that offering mixed contracts (the “mixed” suppliers) is always beneficial to both the customers and the resource suppliers. Also, under price competition, the market will be segmented and both the dynamic and static suppliers may make profit at the equilibrium.

A simple example

Let us suppose, for example, that at the start of a year a customer can sign a contract for a fixed line at a cost of $\$a$ per 1 Mbps. Additionally, he may purchase at the start of each week, additional capacity for a cost of $\$b$ per 1 Mbps per year. Assuming $a < b$, there is the incentive to book ahead. This is easier for customers whose requirements are more predictable. Suppose that the supplier has no capacity constraint. Consider now three customers and their bandwidth requirements. For simplicity we assume that these requirements are inelastic.

Customer 1 needs 1 Mbps pipe throughout the year. The most he would be prepared to pay for this is $\$100$. Customer 2 needs a 1 Mbps pipe on half of the weeks of the year and a 2 Mbps pipe on the other half of the weeks of the year. These weeks are randomly distributed in the year, but are predictable one week ahead. The most he would be prepared to pay for this is $\$190$. Customer 3 needs a 1 Mbps pipe on 9/10ths of the weeks of the year and a 3 Mbps pipe on the other 1/10ths of the weeks of the year, which are again randomly distributed in the year and predictable one week ahead. The most he would be prepared to pay for this is $\$136$. Customers 2 and 3 will only purchase a contract if it can satisfy their requirements on all weeks of the year. The problem for the supplier is to choose a and b so as to solve the problem

$$\text{maximize } \sum_{i=1}^3 y_i,$$

subject to $y_1 = a$ if $a \leq 100$, otherwise 0, $y_2 = \min\{2a, a + b/2\}$ if $\min\{2a, a + b/2\} \leq 190$, otherwise 0, and $y_3 = \min\{3a, a + 2b/10\}$ if $\min\{3a, a + 2b/10\} \leq 136$, otherwise 0.

If we allow only the year-long contracts, (so that effectively $b = \infty$), then by taking $a = 95$ the revenue is maximized to $95 + 190 = \$285$. Only the first two customers purchase contracts. If we have wanted to ensure that all three customers

buy contracts, we would have needed to set $a = 135/3 = 45$ and this would have produced revenue of only $\$270$.

Now suppose we allow bookings to be made just one week ahead. It will be optimal to take $a = 100$ and $b = 180$. All customers will make purchases and the revenue will be $100 + 190 + 136 = \$426$. Note that we are requiring customer 1 to pay $\$10$ more than before. If we wished to avoid this we might take $a = 95$ and $b = 190$. Now we can be certain that customers 1 and 2 will purchase the same as before, and the revenue will be $95 + 190 + 133 = \$418$.

The nice thing about this last solution is that we need not know before introducing the new tariff that the third customer even exists. If he does, then we increase revenue by 46.7%. But if he does not exist then nothing has been lost. If we are lucky, there will be even further customers, additional to customer 3, who will start buying contracts.

This example suggests that a provider may increase its profit by using a mixed contract and that customers may obtain at least the same net benefit they obtained in the case of using optimal static contracts. But what if the provider uses purely dynamic contracts? In this case the optimal revenue is $\$370$, for $b = 100$, which is larger than the optimal static revenue of $\$285$. But this is not always the case! If we have a different percentage of the same customers, the optimal static revenue may be higher. For instance, if no customer 3 is present, then the maximum revenue of static, dynamic and mixed contracts are $\$285$ ($a = 95$), $\$250$ ($b = 100$) and $\$290$ ($a = 100$, $b = 180$) respectively. Hence dynamic contracts may no longer be optimal for the provider.

This paper is organized as follows. In Section II, we formulate the general optimization problem that a customer faces when provided with a static, mixed or dynamic contract. Section III defines the utility function that will be used throughout the rest of the paper and in Section IV, we use this utility function to model the customer’s net benefit maximization problem. In Section V, we model the provider’s revenue maximization problem and we show how the prices of the various contracts are formed. We also consider various cases of customer distributions. Section VI models a price competition game between providers of different type of contracts, where we give an insight of how prices will adjust. One can prove that, under certain conditions, at the equilibrium the market will be segmented between the providers. Section VII concludes our work and provides some points for future work.

II. THE OPTIMIZATION PROBLEM

Let us start from the position that contracts generate charges over fixed long time periods consisting of n shorter periods (slots). Prices are assumed to be fixed and known to customers. They buy long-term contracts where they secure a certain fixed amount of resources for the n slots, and may combine these with short-term contracts for acquiring additional resources in each slot.

For simplicity assume that the long time period is a year and slots correspond to weeks, i.e., $n = 52$. There are N customers who are prepared to buy long-term contracts given that we are charging a static price corresponding to $\$a_0$ per

1 Mbps per week (the actual price is na_0 for the long-term period but we like to express it per slot). Customer i buys bandwidth of x_i Mbps in this long-term static contract. The revenue is presently

$$r(a_0, b = \infty) = na_0 \sum_{i=1}^N x_i.$$

Suppose now that we introduce the possibility of buying further units of bandwidth, one-week-ahead, for $\$b$ per 1 Mbps. Suppose that for the coming week, customer i has a utility for bandwidth that is parameterized by θ , which he can predict one week ahead. If he already owns a year-long contract for bandwidth x , and he additionally buys bandwidth y for the coming week, then his utility is $u_i(\theta, x + y)$. This customer will choose to buy a year-long contract for x_i where this comes from his finding his maximum net benefit as

$$nb_i(a, b) = \max_{x_i} \left\{ nE \left[\max_{y_i} \{u_i(\theta, x_i + y_i) - by_i\} \right] - na x_i \right\}, \quad (1)$$

where the expectation is taken over θ and $n = 52$.

Observe first that the solution of the optimization problem is independent on n when both prices are expressed on a per slot basis. The average revenue he generates is $r_i(a, b) = na x_i + nb \bar{y}_i$, where $\bar{y}_i = E[\arg \max_{y_i} \{u_i(\theta, x_i + y_i) - by_i\}]$ is the average bandwidth he will buy in a week in the dynamic bandwidth market, and x_i is the maximizer in (1). The total revenue may be denoted as $r(a, b) = \sum_i r_i(a, b)$. We seek to maximize $r(a, b)$.

The problem faced by the supplier offering the tariff (a, b) becomes more complex if we consider the possibility that the customer may seek another supplier offering the tariff (a_1, b_1) if his net benefit using this supplier is higher. For instance, suppose the second supplier offers a purely static contract (a_0, ∞) . This suggests for a market consisting of a single customer the optimization problem

$$\underset{a, b}{\text{maximize}} \ r(a, b) \text{ s.t. } \ nb(a, b) \geq nb(a_0, \infty). \quad (2)$$

Intuition suggests that at optimality, $a \leq a_0$ and $b > a$. We also observe that for customers with constant bandwidth requirements introducing the possibility of buying bandwidth on demand does not increase the revenue of the supplier. But in the case of customers with fluctuating requirements, (2) should lead to an increase of revenue.

We can derive now an important property of static versus purely dynamic contracts. Under the same price $a = b$, dynamic contracts are more beneficial for the customer. This is easy to see from (1) since

$$\max_x \left\{ nE \left[u_i(\theta, x) - ax \right] \right\} \leq nE \left[\max_y \{u_i(\theta, y) - ay\} \right].$$

A simpler version of the problem is one in which customers have fluctuating but predictable bandwidth requirements. This corresponds to their knowing the realization of the values of θ for the different weeks in advance. In this case the customer solves

$$nb_i(a, b) = \max_{x_i} \left\{ \sum_{j=1}^n \left[\max_{y_i^j} \{u_i(\theta_j, x_i + y_i^j) - by_i^j\} \right] - na x_i \right\}, \quad (3)$$




Fig. 1. The utility function and resulting demand function when $k = 1$.

where y_i^j is the bandwidth to be bought from the dynamic market during the j th period. Now the revenue generated is $r_i(a, b) = na x_i + b \sum_j^n y_i^j$.

This model allows us to formulate some more interesting problems. Observe that the dynamic demand y_i^j of customer fluctuates over time and its statistics are affected by the choice of a, b . If there is finite capacity in the system that must be procured at the beginning of the year, then we must also make sure that the total demand from customers rarely exceed the capacity. This suggests that we associate an effective bandwidth [7] with each customer i capturing the variability of y_i^j . Lets denote this by $eb_i(a, b)$. Then a related optimization problem is

$$\underset{a, b}{\text{maximize}} \ r(a, b) \text{ s.t. } \sum_i eb_i(a, b) \leq C. \quad (4)$$

This captures the fact that increasing a and decreasing b may result in higher revenues but it creates more fluctuating demand which may be harder to provide.

III. A MODEL FOR CONSUMER UTILITY

We suppose that at the time that a customer purchases static bandwidth he does not yet know his utility function for bandwidth. He knows only that it will be of the form $u_k(x)$, where k is a parameter, presently unknown, but distributed a priori as a random variable with a known distribution function $F(k)$. In particular, we will illustrate ideas with

$$u_k(x) = \begin{cases} kx - \frac{1}{2}x^2, & x \leq k, \\ \frac{1}{2}k^2, & x \geq k. \end{cases} \quad (5)$$

Assuming this utility function, if the user faces a static price of a and knows k then he will choose x to maximize $u_k(x) - ax$. This gives the demand function $x_k(a) = \max\{k - a, 0\}$ (See Fig. 1).

IV. THE CONSUMER'S PROBLEM: MAXIMIZING NET BENEFIT

Consider a single consumer and the purchases of bandwidth that he will make when faced with a static, dynamic, or mixed supplier, and where the static and dynamic bandwidth prices are a and b respectively. If facing a static, dynamic or mixed supplier, the customer optimizes over the quantity of bandwidth he buys and obtains net benefits (average, per slot)

of

$$\begin{aligned} nb_S(a) &= \max_x \left\{ E[u_k(x) - ax] \right\}, \\ nb_D(b) &= E \left[\max_x \{u_k(x) - bx\} \right], \\ nb_M(a, b) &= \max_x \left\{ E \left[\max_y \{u_k(x+y) - by\} - ax \right] \right\}, \end{aligned}$$

where the expected values are taken over k .

To illustrate the sort of things that can occur, we consider examples in which the utility function is of the form (5) and k has some special distributions. One of these is where k is arbitrarily distributed on $[0, 1]$ (see Section V-A). Let us consider for the moment only the cases of static and dynamic sellers (not mixed). We then have,

$$\begin{aligned} nb_S(a) &= \max_x \left\{ \int_0^1 u_k(x) dk - ax \right\}, \\ nb_D(b) &= \int_0^1 \max_x \{u_k(x) - bx\} dk. \end{aligned}$$

Another possible distribution for k is when $k = k_1$ or $k = k_2$ with probabilities $1 - p$ and p respectively (see Section V-B). We suppose $0 \leq k_1 < k_2$. We now have

$$nb_S(a) = \max_x \left[(1-p)u_{k_1}(x) + pu_{k_2}(x) - ax \right].$$

For a fixed a , this is a convex function of p . For dynamic contracts the user's net benefit is

$$nb_D(b) = (1-p) \max_x [u_{k_1}(x) - bx] + p \max_x [u_{k_2}(x) - bx].$$

For a fixed b , this is a linear function of p . Thus, there will be some range $p \in [p_1, p_2] \subset [0, 1]$ for which the dynamic contract will be preferred. To see this, note that if $a = b$ then the user will prefer to buy dynamically, except at $p = 0$ and $p = 1$, where the user is indifferent. As b increases, $nb_D(b)$ decreases at every value of p and remains a linear function of p . Thus as functions of p the linear function $nb_D(b)$ crosses the convex function $nb_S(a)$ at most twice.

V. THE MONOPOLIST SUPPLIER'S PROBLEM

A. Identical customers with arbitrarily distributed k

Suppose there is a single monopolist supplier who can supply bandwidth at cost c and who is attempting to sell to a population of identical customers. Let us ask whether he can make more profit as a static, dynamic or mixed supplier. To discover the answer, we must perform a somewhat complicated calculation. We assume that the supplier knows F , the distribution of k and can determine a typical customer's demand function, say $x_S(a)$ in the static case. He then chooses the price a to maximize the profit $(a - c)x_S(a)$.

The results of Table I are obtained numerically under the assumption that k is uniformly distributed on $[0, 1]$. We suppose $c = 0$ so that the supplier's revenue and profit are the same thing. We show what happens for the three types of supplier in the columns labeled (S), (M) and (D). The figures for revenue and mean bandwidths bought and sold are per customer.

TABLE I

(S) THE OPTIMAL PURELY STATIC CONTRACT HAS $a = 2/9$ AND THE OPTIMAL REVENUE IS $2/27$. THE USER BUYS $x = 1/3$. THE USER'S NET BENEFIT IS $7/162$. THE MEAN BANDWIDTH CONSUMED IS $1/3$. (M) THE OPTIMAL MIXED CONTRACT HAS $a = 6/25$ AND $b = 2/5$ AND OPTIMAL REVENUE $2/25$. THE USER BUYS $x = 1/5$ STATIC AND HIS NET BENEFIT IS $11/250$. THE MEAN BANDWIDTH CONSUMED IS $7/25$. (D) THE OPTIMAL PURELY DYNAMIC CONTRACT HAS $b = 1/3$ AND OPTIMAL REVENUE $2/27$. THE USER'S NET BENEFIT IS $4/81$. THE MEAN BANDWIDTH CONSUMED IS $2/9$.

	(S)	(M)	(D)
seller revenue	0.0741	0.0800	0.0741
optimal a	0.2222	0.2400	
optimal b		0.4000	0.3333
bandwidth bought in static	0.3333	0.2000	
mean bandwidth sold	0.3333	0.2800	0.2222
user net benefit	0.0432	0.0440	0.0494

These numerical results suggest some conjectures. If true in general, they support the notion that there is advantage in offering mixed contracts.

- The revenue achieved by the seller is strictly greater in (M) than in (S) or (D).
- The mean bandwidth sold decreases from (S) to (M) to (D).
- The user's average net benefit increases from (S) to (M) to (D).

We also notice the interesting fact that the revenue achieved by the seller is the same in (S) and (D). This turns out to be true in general, assuming the utility is of form (5). We state this as follows.

Proposition 1: Suppose the utility function is of the form in (5). Then a seller of static contracts who maximizes his revenue by choice of a obtains the same revenue as a seller of dynamic contracts who maximizes his revenue by choice of b . Moreover, at these optimums the bandwidth x^* that is sold by the seller of static contracts is equal to the optimum dynamic contract price b^* .

Proof: The theorem holds whatever the distribution of k , but let us suppose for simplicity that k is arbitrarily distributed over $[0, 1]$ with density function $f(k)$. If the seller chooses a price a for static contracts, the buyer will choose x to maximize

$$\int_0^x u_k(k) f(k) dk + \int_x^1 u_k(x) f(k) dk - ax.$$

This is maximized where

$$-a + \int_x^1 (k - x) f(k) dk = 0.$$

Thus the seller maximizes his revenue by maximizing over x

$$x \int_x^1 (k - x) f(k) dk. \quad (6)$$

Now if the price of dynamic contracts is b , then the amount of bandwidth purchased is $(k - b)^+$. Thus the seller seeks to maximize over b

$$b \int_b^1 (k - b)f(k) dk. \quad (7)$$

The theorem follows by comparing (6) and (7). ■

Note also that the optimal x is where

$$\frac{d}{dx} \left\{ x \int_x^1 (k - x)f(k) dk \right\} = \int_x^1 (k - 2x)f(k) dk = 0.$$

So taking $x = x^*$ as the solution,

$$\int_{x^*}^1 (k - x^*)f(k) dk = (1 - F(x^*))x^* < x^*.$$

Thus as $x^* = b^*$ we have that the bandwidth sold under the dynamic contract is less than the bandwidth sold under the static contract, as in Conjecture B.

B. Identical customers with two-point distributed k

Since we would like to consider what happens when customers are not identical we wish to find some simple model in which their differences are captured by a single parameter. So let us suppose that $k = k_1$ or $k = k_2$ with probabilities $1 - p$ and p . Now different values of p can distinguish customers from one another. Initially, let us fix p and consider what happens when a monopolist is selling to a population of customers, all of whom have the same p .

1) *Static contracts*: The customer's net benefit is

$$p_1 u_{k_1}(x) + p_2 u_{k_2}(x) - ax.$$

This is a continuous function of x , and it has a continuous first derivative. Taking the first derivative, we see that the price at which the user will purchase x is

$$p_1 u'_{k_1}(x) + p_2 u'_{k_2}(x) = p_1(k_1 - x)^+ + p_2(k_2 - x)^+.$$

In fact, we can write

$$p_1(k_1 - x)^+ + p_2(k_2 - x)^+ = \max\{0, \bar{k} - x, p_2(k_2 - x)\}.$$

Or we can also invert this to write the demand function as

$$x(a) = \max\{0, \bar{k} - a, k_2 - a/p_2\}.$$

Thus the seller's maximum revenue is

$$\begin{aligned} & \max_x x \max\{0, \bar{k} - x, p_2(k_2 - x)\} \\ &= \max\left\{0, \max_x \{x(\bar{k} - x)\}, p_2 \max_x \{x(k_2 - x)\}\right\} \\ &= \max\left\{\frac{1}{4}\bar{k}^2, \frac{1}{4}p_2 k_2^2\right\} \end{aligned}$$

Now $\frac{1}{4}\bar{k}^2 \geq \frac{1}{4}p_2 k_2^2$ if and only if $p_2 \leq p^*$ where

$$p^* = \frac{1}{(k_2/k_1 - 1)^2},$$

since $p_1 = 1 - p_2$. This gives the following.

Proposition 2: The optimal (a^*, x^*) are either $(\frac{1}{2}\bar{k}, \frac{1}{2}\bar{k})$ or $(\frac{1}{2}p_2 k_2, \frac{1}{2}k_2)$ as p_2 is less or greater than p^* .

Note that $2k_1 \geq k_2 \implies p^* \geq 1 \implies p_2 \leq p^*$ and the optimum is at $x^* = \frac{1}{2}\bar{k}$ for all values of p_2 . Conversely, $p_2 > p^* \implies 2k_1 < k_2$.

2) *Dynamic contracts*: If the seller sets price b then the user will buy average bandwidth of

$$p_1(k_1 - b)^+ + p_2(k_2 - b)^+.$$

We know by Proposition 1 that the optimal value of b is $\frac{1}{2}\bar{k}$ or $\frac{1}{2}k_2$. If $p_2 \leq p^*$ we have $b = \frac{1}{2}\bar{k}$. If $p_2 \geq p^*$ we have $b = \frac{1}{2}k_2 > k_1$. It turns out that in both cases the average bandwidth is the same as in the case of the static contract. Hence, Conjecture B is not true in a strict sense for the case of the two-point distributed k . But under a arbitrarily distributed k , Conjecture B is true in a strict sense (as shown in Section V-A).

Furthermore, we can compute the difference in the user's net benefit under optimal purely dynamic and purely static contracts as

$$nb_D - nb_S = \begin{cases} \frac{1}{2}p_1 p_2 (k_2 - k_1)^2, & p_2 \leq p^* \\ -\frac{1}{2}p_1 k_1^2, & p_2 > p^* \end{cases}$$

This gives

Proposition 3: If $p_2 > p^*$ the a user strictly prefers to face a monopolist seller of static contracts rather than a monopolist supplier of dynamic contracts. If $p_2 < p^*$ his preference is reversed.

Thus, we see that whereas the optimal purely static and purely dynamic contracts produce the same revenue for the seller, the type of contract that obtains the greater net benefit for the user depends on whether p_2 is less or greater than p^* . Notice that there is a discontinuity in the user's net benefit at $p_2 = p^*$. Conjecture C is not true in general. It is only true if $p_2 \leq p^*$.

Remark: The above proposition can be generalized, in the sense that for every price a of the static provider, a dynamic provider can publish a price b that offers him at least the same revenue, while the customer acquires a higher net benefit, if $p_2 \leq p^*$. The result is reversed in the case of $p_2 > p^*$. In Fig. 2 we show how the user's net benefit can vary with the seller's revenue as the seller varies his static price a , or dynamic price b . The sellers have no cost. Notice that the maximum revenue that can be obtained by either method is the same.

3) *Mixed contracts*: Consider the optimal mixed strategy, optimized over prices $a \leq b$. Suppose the net benefit at the optimum is

$$p_1 u_{k_1}(x + y_1) + p_2 u_{k_2}(x + y_2) - xa - p_1 y_1 b - p_2 y_2 b.$$

Note first that we cannot have $y_1 > 0$, since as $b > a$ the net benefit to the user could be improved by taking $x \rightarrow x + y_1$, $y_1 \rightarrow 0$ and $y_2 \rightarrow y_2 - y_1$. Therefore we may consider the problem when the net benefit is of the form

$$p_1 u_{k_1}(x) + p_2 u_{k_2}(x + y_2) - xa - p_2 y_2 b.$$

Now we must have $p_2 b < a$, otherwise the user does best to buy only static. Now also note that we cannot have $x > k_1$. For if this were so, the user could make the change $x \rightarrow x - \epsilon > k_1$ and $y_2 \rightarrow y_2 + \epsilon$, and his net benefit would increase by $(a - p_2 b)\epsilon > 0$. Thus the optimum must occur where $x \leq k_1$. Note also, that the optimum must occur where $x + y_2 \leq k_2$.

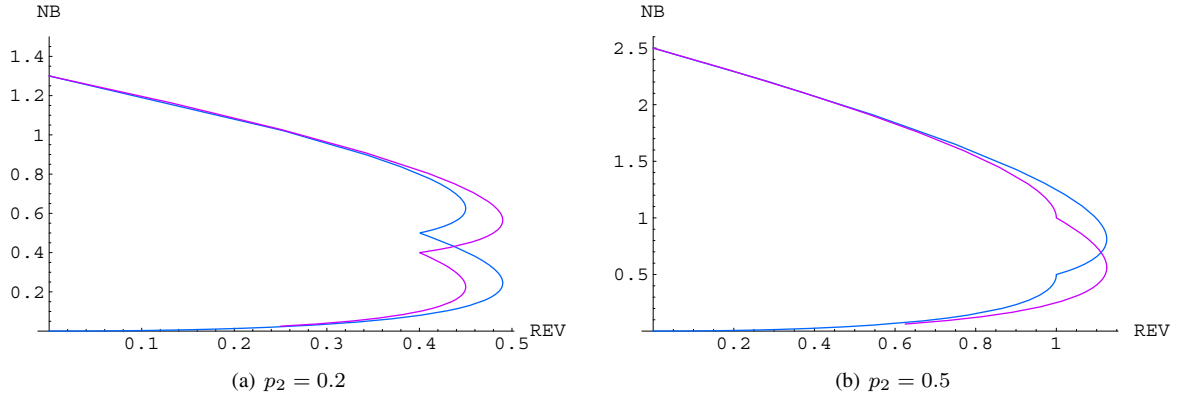


Fig. 2. The user's net benefit (NB) against the seller's revenue (REV). The static seller is shown in blue and the dynamic seller in violet. In this example, $k_1 = 1$, $k_2 = 3$. If $p_2 < p^* = 0.25$, the seller of dynamic (violet) can, at any revenue value, offer the user a greater net benefit than can the seller of static (blue). The reverse is true if $p_2 > p^*$.

For a given x , such that $x < k_2 - b$, the user optimally takes $y_2 = (k_2 - x - b)$, and so his net benefit is then

$$p_1 u_{k_1}(x) + p_2 u_{k_2}(k_2 - b) - xa - p_2(k_2 - x - b)b.$$

The derivative of this with respect to x is

$$-a + p_2 b + p_1(k_1 - x).$$

Since $-a + p_2 b < 0$, this is negative at $x = k_1$, and so the optimum is at some $x < k_1$, where $x = k_1 - b + (b - a)/p_1$. Thus the seller's revenue can be found by substituting this value of x into his revenue of $xa + p_2(k_2 - x - b)b$ and optimizing over a and b . This gives

$$a^* = \frac{1}{2}\bar{k} \quad \text{and} \quad b^* = \frac{1}{2}k_2.$$

Note that these satisfy $b > a > p_2 b$, as required. The optimal purchases are $x = \frac{1}{2}k_1$ and $y_2 = \frac{1}{2}(k_2 - k_1)$. Thus the mean amount that is purchased is $x + p_2 y_2 = \frac{1}{2}\bar{k}$. Interestingly, we have the same static price, and same average bandwidth purchased, as in the static contracts case, provided we are in the regime where $p_2 \leq p^*$. As commented previously, a special case is $2k_1 \geq k_2$, which implies $p^* = 0$.

The user's net benefit can be calculated to be

$$p_1 u_{k_1}(k_1/2) + p_2 u_{k_2}(k_2/2) - (1/4)(p_1 k_1^2 + p_2 k_2^2) = \frac{1}{8}(p_1 k_1^2 + p_2 k_2^2)$$

Under the static contract it was

$$\frac{1}{8}(p_1 k_1 + p_2 k_2)^2 = \frac{1}{8}\bar{k}^2,$$

which is less. The seller's revenue has increased from $\frac{1}{4}\bar{k}^2$ to $\frac{1}{4}(p_1 k_1^2 + p_2 k_2^2)$. It is interesting to note that both the seller and user increase their profit and net benefit, respectively, by the same factor of

$$\frac{p_1 k_1^2 + p_2 k_2^2}{(p_1 k_1 + p_2 k_2)^2}.$$

The greatest possible value for this ratio in the region $2k_1 \geq k_2$ is 1.125.

TABLE II
SUMMARY OF RESULTS

$p_2 \leq p^*$	$p_2 > p^*$
$r_S = r_D < r_M$	$r_S = r_D < r_M$
$nb_S < nb_M < nb_D$	$nb_D < nb_S < nb_M$
$m_S = m_M = m_D$	$m_S = m_D > m_M$

Remark: Note that the seller obtains the same revenue as he could obtain if he knew k and used 'personalized pricing', i.e., charging a price $\rho = \frac{1}{2}k_1$ if $k = k_1$, and $\rho = \frac{1}{2}k_2$ if $k = k_2$.

Let us summarize results. In Table II, we write the seller's revenue as r_S, r_M, r_D , the user net benefits as nb_S, nb_M, nb_D and the mean bandwidths sold as m_S, m_M, m_D .

These results suggest that mixed contracts are always preferable to static contracts, in the sense that $r_M > r_S$, $nb_M > nb_S$ and $m_M \leq m_S$.

VI. PRICE COMPETITION AMONGST SUPPLIERS

We now turn to investigate what happens when suppliers compete with one another on price. Who wins? Is it the supplier of static, dynamic or mixed contracts? Is there an equilibrium in which suppliers of different types can both make positive profit?

Let us suppose that for a given customer $k = 0$ or $k = k_2$, with probabilities $1 - p$ and p . We suppose $p \in [0, 1]$ is distributed across the population of users with a density function of $f(p)$.

As previously, the static provider sells at price a , but the contract for purchase must be made before the customer knows whether his k is equal to 0 or 1. The dynamic provider sells at price b , and with the flexibility that the customer need not make the purchase until he knows whether his k is equal to 0 or 1.

If buying from the static provider the customer chooses x to maximize

$$pu(1, x) - ax,$$




Fig. 3. The reaction curves $a(b)$ against b (in blue), and a against $b(a)$ (in violet), for different distributions of p .

and so optimally buys $\max\{1 - a/p, 0\}$. If buying from the dynamic provider the customer chooses x to maximize

$$pu(1, x) - pbx,$$

and optimally buys $\max\{1 - b, 0\}$. Thus a customer strictly prefers to buy from static rather than dynamic if and only if $a < pb$.

Suppose the two providers have unit cost of c_1 and c_2 , respectively. Then the profits obtained by the two providers are

$$f_A(a, b) = \begin{cases} (a - c_1) \int_{a/b}^1 (1 - a/p) f(p) dp, & a < b \\ 0, & a \geq b \end{cases}$$

$$f_B(a, b) = \begin{cases} (b - c_2) \int_0^{a/b} p(k - b) f(p) dp, & a < b \\ (b - c_2) \int_0^1 p(k - b) f(p) dp, & a \geq b \end{cases}$$

These are the payoffs in a Bertrand game of price competition. From these we can compute the reaction curves.

$$a(b) = \arg \max_a \{f_A(a, b)\}, \quad b(a) = \arg \max_b \{f_B(a, b)\}.$$

To give some numerical examples, suppose $c_1 = 0.1$ and $c_2 = 0.2$. Fig. 3(a) shows the reaction curves when p is uniformly distributed, i.e., $f(p) = 1$. Fig. 3(b) shows these curves when $f(p) = 6p(1 - p)$. Fig. 3(c) shows these curves when the distribution is more concentrated around $p = 1/2$, with $f(p) = 630p^4(1 - p)^4$. The point of intersection in each graph is a Nash equilibrium. In the equilibrium point of Fig. 3(a), $a = 0.2105$, $b = 0.3333$ and the respective revenues are 0.0300 and 0.0177.

As already mentioned in Section IV, at the equilibrium point the market will be segmented between the two providers. The customers that reside in the edges of the p 's distribution will prefer the static provider and the rest customers will prefer the dynamic provider.

VII. CONCLUSIONS & FUTURE WORK

We have presented a model for analyzing the different types of a contract in a market of resources that can be provisioned on demand. Our aim was to show that, when using the appropriate pricing scheme, the provider's goal of revenue maximization is not opposed to the consumer's goal for net benefit maximization. We have captured the fluctuating demand of the customers with a single parameter (p), leading to a formulation of high and low demand periods. Based on this model, we have studied the characteristics of such

contracts, under specific distributions of k . We have seen how the provider's revenue, the customer's net benefit and the mean bandwidth sold vary from one type of contract to another and we have provided the conditions under which a mixed contract is better from the rest for both the provider and the consumer. We have also proved that the optimal revenue of a static provider is the same with the revenue achieved by a dynamic provider, under any distribution of k . Finally, we have provided some numerical results for a price competition game between a static and a dynamic provider that shows that a Nash equilibrium point exists and that the market at this point is segmented, hence both providers make profits.

Many issues are open for further research. An interesting point is to see what segment of the market is obtained by a mixed provider when participating in a price competition game with a static or dynamic provider. Furthermore, apart from the example with the two-point distribution of k , results from more generic distributions will be studied. The case of non-identical customers needs also to be studied. Extensions with effective bandwidth will also be considered. An important question that may rise is how the shape of the utility function affects all the aforementioned results. Finally, an open issue is how such a model can be extended for Grid resources, since the definition of the basic resource in Grids and its characteristics are not yet well-defined.

REFERENCES

- [1] "DSL Forum's BroadbandSuite: Unleashing the Power of Broadband," Technical Paper, DSLForum.org.
- [2] "Creating a Commercially Sustainable Framework for IP Services Realizing Next Generation Revenues," White paper, IPSphere Forum, May 2006.
- [3] "GRID COMPUTING on SUN: Deploying Scalable and Efficient HPC Grids in the Modern Commercial or Technical Enterprise," White Paper, SUN Microsystems Inc., March 2006.
- [4] B.M. Mitchell and I. Vogelsang, "Telecommunications Pricing Theory and Practice," *Cambridge University Press*, 1991.
- [5] T. Bergstrom and J.K. MacKie-Mason, "Some Simple Analytics of Peak-Load Pricing," *The RAND Journal of Economics*, vol. 22(2), pp. 241-249, 1991.
- [6] W.Y. Oi, "A Disneyland Dilemma: Two-part Tariffs for a Mickey Mouse Monopoly," *Quarterly Journal of Economics*, vol. 85, pp. 77-96, Feb. 1971.
- [7] F.P. Kelly, "Notes on effective bandwidth," In: F. Kelly, S. Zachary and I. Ziedins (Eds.), "Stochastic Networks: Theory and Applications in Telecommunication Networks," Volume 4 of *Royal Statistical Society Lecture Notes Series*, Oxford, pp. 141-168, Oxford University Press, 1996.