Journal of Universal Computer Science, vol. 14, no. 5 (2008), 766-785
submitted: 1/11/07, accepted: 28/2/08, appeared: 1/3/08 © J.UCS

Dynamic Bandwidth Pricing: Provision Cost, Market Size, Effective Bandwidths and Price Games

Sergios Soursos, Costas Courcoubetis
(Athens University of Economics and Business, Greece
{sns, courcou}@aueb.gr)

Richard Weber
(University of Cambridge, UK
rrw1@statslab.cam.ac.uk)

Abstract: Nowadays, in the markets of broadband access services, traditional contracts are of "static" type. Customers buy the right to use a specific amount of resources for a specific period of time. On the other hand, modern services and applications render the demand for bandwidth highly variable and bursty. New types of contracts emerge ("dynamic contracts") which allow customers to dynamically adjust their bandwidth demand. In such an environment, we study the case of a price competition situation between two providers of static and dynamic contracts. We investigate the resulting reaction curves, search for the existence of an equilibrium point and examine if and how the market is segmented between the two providers. Our first model considers simple, constant provision costs. We then extend the model to include costs that depend on the multiplexing capabilities that the contracts offer to the providers, taking into consideration the size of the market. We base our analysis on the theory of effective bandwidths and investigate the new conditions that allow the provider of dynamic contracts to enter the market.
Key Words: network economics, contracts, pricing, bandwidth on demand, effective bandwidths, statistical multiplexing, reaction curves, competition, provision cost.
Category: C.2.0, G.1.6, J.4

1 Introduction

In the bandwidth markets, contracts between providers and customers have traditionally been of static nature. Both parties agree on the terms of a long-term contract, typically of one year, and comply to their contractual obligations throughout the defined time period. Such *static contracts* are widely used in the broadband access markets and for the formation of VPNs. The customers of the aforementioned services vary from small home users to large corporate users. Such contracts have the form of "buy X amount of bandwidth at a total price of \$Y, for one year" or, a more generic form of "buy any amount of bandwidth at a unit price of \$Y, for one year".

On the other hand, the nature of demand for such services cannot be characterized as 'static'. Especially for large customers (companies, organizations, etc), the demand is highly variable, depending on the type of applications/services that run above these lower-level (communication) services. Furthermore, new technologies allow for dynamic allocation of network resources, based on the customer's current demand for bandwidth [Rabbat and Hamada 2006]. This change of paradigm renders *dynamic contracts* more suitable. A dynamic contract is also valid for a long-term period (e.g. for one year), but offers the customer the ability to buy bandwidth on shorter timescales (e.g. per week), according to his current needs and at a price, published by the provider, that remains constant throughout the long-term period, e.g. "buy any amount of bandwidth each week at a unit price of $Z, for one year".

The above contracts do not apply only to the case of a bandwidth market. Consider, for example, the case of a commercial Grid platform, a Grid marketplace, where providers and customers meet to buy and sell Grid resources. A static or dynamic contract in such a context would involve the purchase, or better leasing, of a number of virtual machines (VMs) for a specific time period. In such a market, a broker could offer longer term contracts, guaranteeing a fixed price and offering two alternative ways of acquiring resources: a fixed amount of VMs for a long period of time (static contract) or a more flexible contract for a variable amount of resources (dynamic contract) over a long period.

The basic properties of static and dynamic contracts in a monopoly environment have have been studied in our previous work [Courcoubetis et al. 2006]. In this paper, we examine the properties of such contracts under a price competition setup. Assume that there are only two providers, one that offers static contracts (*static provider*) and one offering dynamic contracts (*dynamic provider*). The providers must choose the unit price to publish in order to maximize their profits. Each provider takes into consideration the price that his competitor has published, in order to publish his own, profit-maximizing, unit price. Hence, a price competition situation emerges, where each provider is able to calculate a 'response' for every possible action his competitor takes, assuming the complete information case.

We also assume that both types of contract have the same duration. Thus, at the beginning of the long-term period, the customer must choose between a static and a dynamic contract, taking into consideration the unit prices that both contracts publish, as well as his (expected) requirements in bandwidth, in order to maximize his net benefit.

The most important factor that affects the provider's profit maximization problem is the provision cost. This cost defines what the unit price should be in order for the provider to be able to cover his costs and make some profits. We consider the case that both providers buy wholesale capacity (or dimension their networks) at the beginning of the long-term period, prior to fulfilling their contractual agreements. The provider of static contracts needs only to procure the exact amount of capacity he has sold in his contracts, since his customers cannot change their demand for the duration of the contract. On the contrary, the provider of dynamic contracts has to procure more capacity than the expected average requested in each short-term period, so as to be able to fulfill his contractual obligations. So, it is obvious that the provision cost of a unit of bandwidth for a static provider is less than the unit cost of the dynamic provider.

In our work, we consider two cases as far as the provision cost is concerned. Initially, we provide a summary of the case where the provision costs are constant and independent of the market size. We then consider the case where the provision cost of the dynamic provider depends on the market size and the multiplexing opportunities that dynamic contracts offer. According to the effective bandwidths theory and due to the variable nature of demand under a dynamic contract, the amount of bandwidth reserved per customer in the bandwidth inventory tends to the average bandwidth consumption of the contract as the size of the market increases, i.e. as the number of dynamic contracts served increases. Hence, our purpose is to capture the effects of statistical multiplexing in the cost structure of the dynamic provider.

The main results of this paper involve the operating point of the market. In other words, for both cases of constant and multiplexing-dependent costs, we examine if and when there exists an equilibrium in the price competition game and how the market is segmented at this point. Especially for the more interesting case of the multiplexing-dependent provision costs, we examine what is the necessary size of the market for a dynamic provider to enter and become active as the provision costs increase, what are the profits of the providers as the market size increases and how the amount of the bandwidth reserved per customer changes with the market population.

Competition between providers has also been studied in [Mason 2000], where it is shown that flat rate pricing is more efficient than usage-based prices in equilibrium, under certain conditions. Note that our pricing model is related to the above model since providers compete by changing their pricing schemes for usage: in our model the customers decide on the size of their contracts both in the static and in the dynamic case. But it is also remotely related to usage based pricing, since customers are not charged for their actual usage. The form of contracts described earlier can be more of a flat rate charging, although the resource consumption by the costumers is not unrestricted. Competition between providers has been studied in relation with service compatibility as well. [Foros and Hansen 2001] model a two-stage game in a duopoly setup, where the providers first choose the level of service compatibility and then decide on prices. They show that increased compatibility reduces competitive pressure due to network externalities. [Gibbens et al. 2000] have shown that 'multiproduct' outcomes of a competition between providers always offer lower profits to them, as opposed to the case of a single service class. In our context, providers offer compatible services, i.e. a single service class, but compete in order to gain as many customers as possible, by offering different pricing schemes for usage. In [Davies et al. 2004], the growing necessity for capacity planning and pricing in packet switched networks is discussed. The relation between statistical multiplexing and effective bandwidths is explained in [Courcoubetis and Weber 2003].

The remainder of this paper is organized as follows. Section 2 presents the main assumptions and defines the demand model. In Section 3, a price competition game is analyzed, where the players (the providers of static and dynamic contracts) have constant provision costs. We define the profit maximization problems and provide the equations for the reaction curves. In Section 4, the game setup is altered in order for the provision costs of the dynamic provider to depend on multiplexing. Effective bandwidths are explained and introduced to our model in Section 4.1. Section 5 shows the main numerical results from the analysis of the model. In Section 5.1 the reaction curves are plotted and the equilibrium points are examined, for various market sizes and unit costs. In Section 5.2, the profits are plotted while in Section 5.3 the allocated bandwidth per customer is examined. Section 5.4 discusses about the market segmentation. Finally, we conclude in Section 6 and discuss issues for further study.

2 Demand model and assumptions

Assume that time is slotted, starting from slot 0 up to slot m. The customer decides what type of contract to make just before slot 0. In the case of static contracts, the amount of the bandwidth purchased will be chosen at that time and will remain constant for the remaining m slots. In the case of dynamic contracts, the customer will choose the bandwidth to be purchased at the beginning of each slot i.

To capture the interesting phenomena in this market, we use a very simple form of linear utility function, with slope -1, which varies randomly between slots. The customer's utility for consuming x units of bandwidth in slot i is $u_k(x)$, where k is a parameter, presently unknown, but distributed a priori as a random variable with a known distribution function $F(k)$. To obtain the simple form of demand mentioned before, we assume that the utility function has the following form:

$$u_k(x) = \begin{cases} kx - \frac{1}{2}x^2 \ , & \text{if } x \le k \\ \frac{1}{2}k^2 & , \text{if } x \ge k \, . \end{cases} \tag{1}$$

This utility function is, by design, increasing and concave, as the first segment of (1) suggests. The second segment implies that utility reaches a maximum and then continues to stay there, even if the the consumption increases. If the second term did not exist, the utility would be decreasing, for $x \ge k$, which is not a desired property in our case.

We assume that, at the beginning of each slot, the value of k becomes known to the customer. If the customer faces a price of a and knows k then he will choose x to maximize $E[mu_k(x) - ax] = mE[u_k(x) - \hat{a}x]$, where $\hat{a} = a/m$ is the static price scaled per slot. For simplifying notation, we use this scaled price as the price a of bandwidth under a static contract. Similarly, without loss of generality, we can assume $m = 1$, since the form of the optimization problem remains the same. This gives the demand function $x_k(a) = \max\{k - a, 0\}$ (see Fig. 1).

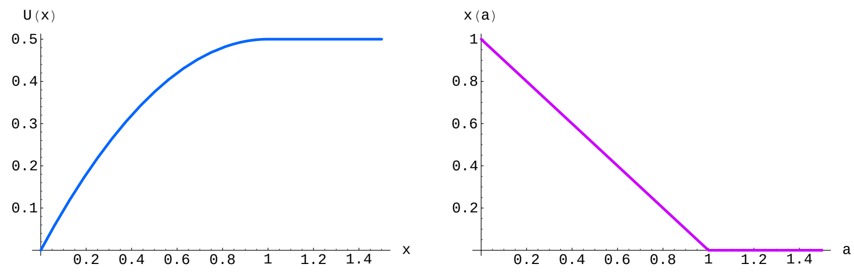

Figure 1: The utility function and resulting demand function when $k = 1$.

In order to illustrate the various issues to be studied and capture the varying bandwidth needs of a customer, we examine a simpler case where k takes discrete values and, more specifically, it is a random variable that follows a two-point distribution. That is, with probability $1 - p$ his utility will be $u_{k_1}(x)$ and with probability p the customer will have a utility of $u_{k_2}(x)$. We suppose $p \in [0, 1]$ is distributed across the population of n customers with a density function of $f(p)$.

The two-point distributed k approach is a simplification of customer's actual behavior regarding network access services. At the same time, it is more realistic than having a constant demand through time. As already mentioned, current applications and services have different requirements in bandwidth. As a result, customers have bursty demand, depending on the applications and services they use each time. In fact, our simplification models a situation where the customer has a minimum and a peak requirement for access bandwidth, without knowing a priori when such needs will occur.

To make things more simple for the analysis to follow, we assume that for any given customer it holds that $k = 0$ or $k = 1$, with probabilities $1 - p$ and p respectively. As already mentioned, the provider of static contracts sells at price a, but the contract must be made prior to the customer's knowledge of the exact allocation of high and low demand periods in the long-term period, i.e. the distribution of k. The dynamic provider sells at price b, with the flexibility that the customer needs not to make the purchase until the beginning of the next slot, when he knows whether his k will be equal to 0 or 1.

When making static contracts, the customer chooses x to maximize $pu_1(x) - ax$ and so optimally buys $\max\{1-a/p, 0\}$. When making dynamic contracts, the customer chooses x to maximize $pu_1(x) - pbx$ and optimally buys $\max\{1-b, 0\}$. Thus a customer strictly prefers a static contract rather than a dynamic one if and only if $a < pb$.

This simplification on the discrete values that k can take (zero or one), does not render the results unrealistic, as related to the more generic case of k_1 and k_2. In fact, in the more generic case, the customer will surely buy k_1 units of bandwidth and the actual decision is how much above the k_1 units will he buy (under either type of contract). Hence, the extra bandwidth will belong to $[0, k_2 - k_1]$. Our only simplification is that we use 1 as an upper bound of the difference.

3 Market equilibrium: Constant cost model

Considering the previous condition for the choice between the two types of contracts, we assume that each provider has a different but constant provision cost. More specifically, the static provider has a unit cost of c_1 and the dynamic provider a unit cost of c_2 which is constant and insensitive to the fact that the provider must provide a total amount of bandwidth that is fluctuating between slots. As explained in the introduction, we assume that it holds $c_2 > c_1$.

Hence, the average profits obtained by the two providers per slot are

$$\mathrm{prof}_S(a,b) = \begin{cases} n(a-c_1)\int_{a/b}^{1}(1-a/p)f(p)\,dp\,, & \text{if } a < b \\ 0 & \text{, if } a \geq b\,, \end{cases} \tag{2}$$

$$\mathrm{prof}_D(a,b) = \begin{cases} n(b-c_2)\int_{0}^{a/b} p(1-b)f(p)\,dp\,, & \text{if } a < b \\ n(b-c_2)\int_{0}^{1} p(1-b)f(p)\,dp & \text{, if } a \geq b\,, \end{cases} \tag{3}$$

where n is the number of customers in the market. Note that n does not affect the profit maximization problem, hence the result is the same even if there was even only one customer.

Equations (2) and (3) provide the payoffs in a Bertrand game of price competition [Binmore 1991]. From these, we can compute the reaction curves:

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

To give some numerical examples, suppose $c_1 = 0.1$ and $c_2 = 0.2$. Figure 2(a) shows the reaction curves when p is uniformly distributed, i.e., $f(p) = 1$. Figure 2(b) shows these curves when $f(p) = 6p(1-p)$. Figure 2(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

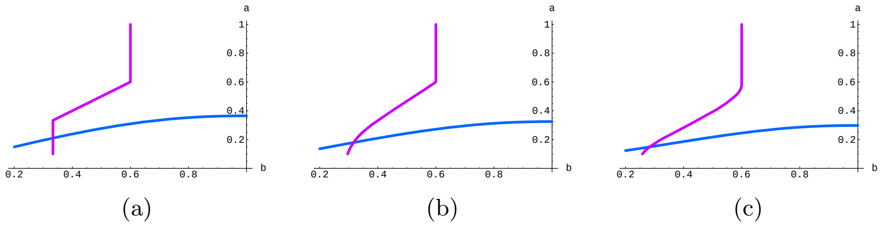

Figure 2: The reaction curves $a(b)$ against b (in blue), and a against $b(a)$ (in violet), for different distributions of p.

the equilibrium point of Fig. 2(a), $a = 0.2105$, $b = 0.3333$ and the respective revenues are 0.0300 and 0.0177.

For the case of a uniformly distributed p, we can provide the function of the $b(a)$ curve and the maximization problem from which we obtain the function of the reaction curve $a(b)$. More specifically, having $f(p) = 1$ and taking into consideration (3), the provider of dynamic contracts expects that the average amount of bandwidth bought by each customer who prefers a dynamic contract is

$$x_D(a,b) = \begin{cases} \frac{1}{2}\frac{a^2}{b^2}(1-b) \ , & \text{if } a < b \\ \frac{1}{2}(1-b) \ , & \text{if } a \geq b \, . \end{cases}$$

Hence, solving the revenue maximization problem for the provider of dynamic contracts

$$\underset{b}{\text{maximize}}(b - c_2)x_D(a,b) \quad \text{w.r.t. } c_2 \leq b \, ,$$

we have that

$$b(a) = \begin{cases} \frac{2c_2}{1+c_2} \ , & \text{if } a < \frac{2c_2}{1+c_2} \\ \frac{1+c_2}{2} \ , & \text{if } a > \frac{1+c_2}{2} \, . \end{cases} \tag{4}$$

The only case that remains to be examined is what value b takes when $\frac{1+c_2}{2} \geq a \geq \frac{2c_2}{1+c_2}$. In this range, for every $b > a$ the profits increase as b decreases. When b becomes lower than a $(b < a)$, the profit maximization problem changes and profits increase as b increases. Hence, the profit maximizing value of b, in the interval $[\frac{1+c_2}{2}, \frac{2c_2}{1+c_2}]$, is where $b = a$.

From (2) we get that the amount of bandwidth bought by each customer that prefers a static contract is

$$x_S(a,b) = \begin{cases} 1 - \frac{a}{b} + a \ln \frac{a}{b} \ , & \text{if } a < b \\ 0 \ , & \text{if } a \geq b \, . \end{cases}$$

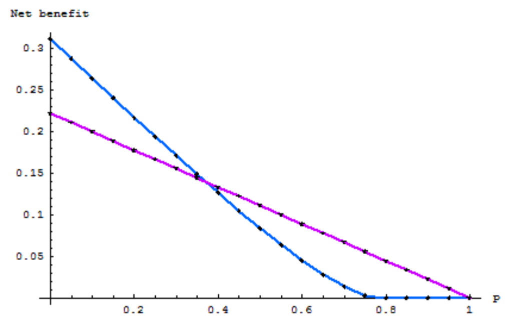

Figure 3: The customer's net benefit for various values of p, with $c_1 = 0.1$ and $c_2 = 0.2$. The violet line depicts the net benefits under a dynamic contract while the blue line shows the net benefits under a static contract. We observe that both providers gain a share of the market, with the dynamic provider acquiring the customers with a high value of p.

Hence, the provider of static contracts wants to maximize his profits, i.e. chooses a price $a(b)$ which solves

$$\underset{a}{\text{maximize}}(a - c_1)x_S(a, b) \quad \text{w.r.t. } c_1 \leq a\,.$$

The exact formula for $a(b)$ cannot be deduced by taking the partial derivative w.r.t a of the above maximizing function equal to zero, due to the complexity of the problem. But the numerical results reveal the form of the curve and the existence of an equilibrium point. As far as the way the market is segmented, Fig. 3 shows how the net benefits of a customer vary with his type p, under a static and a dynamic contract. Hence, the way the market is segmented is obvious: for low to medium values of p, the customers prefer the provider of static contracts over the provider of dynamic contracts. This preference is reversed for higher values of p. The fact that the market is segmented, opposes to the classic theory of Bertrand games where one provider (the one with the lower costs) attracts all the customers.

4 Market equilibrium: Multiplexing-dependent cost model

In the previous section, we have considered the case where the unit provision costs under static or dynamic contracts are constant. As already mentioned in the introduction, c_1 and c_2 may not be arbitrary. The size of the customer base and the provider's multiplexing capabilities affect the provision costs, especially for the provider of dynamic contracts.

It is reasonable to assume that under static contracts, the provider reserves the total amount of bandwidth purchased by the customers. Hence, it is relatively easy for a provider of such contracts to plan the capacity of his network, based on the number of his customers and their contracted capacity. On the other hand, under dynamic contracts, users express demand that varies with time. This results in a variable demand for the provider to serve over a long period. Hence, capacity planning is no longer obvious, since the dynamic provider must estimate how much bandwidth to reserve per dynamic contract. Reserving the maximum that a customer may request, may lead to a waste of bandwidth, since it is highly improbable that all customers will simultaneously request this maximum value. If he uses overbooking, then he reduces his provision cost but may incur a cost of not fulfilling his contractual commitments.

It follows that the amount of bandwidth the dynamic provider needs to reserve per contract is higher than the average consumed by the contract. This amount corresponds to the effective bandwidth of the bandwidth demand process generated by the contract over time. We discuss this in the next section.

Suppose now that both providers must define a bandwidth inventory for satisfying the needs of their contracts and that this inventory will be of some fixed size C, at a unit cost c, same for both providers. In the case of supporting static contracts, the size of the inventory will be equal to the total requested demand of the signed contracts. In the case of dynamic contracts, the amount reserved per customer will need to be larger than the average bandwidth that the provider will need to offer per slot to that customer.

It is hence reasonable to assume that in the case of dynamic contracts, the unit cost is directly affected by the number of customers a provider serves. That is, a higher number of customers will lead to lower provision costs since their demand can be better multiplexed by the inventory than in the case of fewer customers. Ideally one could say that, for infinite number of customers, the provision unit cost of such a provider is the same with the provision unit cost of a static contract, since he needs to reserve an amount close to the average amount of bandwidth per contract. The above statement is based on the fact that the more traffic flows are, the more efficient multiplexing becomes.

4.1 Effective bandwidth analysis

Our purpose is to include the relationship between costs and market size into the model of the price game between a static and a dynamic provider. In fact, our goal is to include the notion of effective bandwidth in the definition of the provision cost for the provider that offers dynamic contracts.

In order to do so, we denote with c the cost for providing one unit of bandwidth in the bandwidth inventory. Since the provider of static contracts needs to maintain an inventory equal to the sum of the bandwidth sold in the contracts, the cost of a unit offered is exactly the same with the unit cost of the bandwidth inventory ($c_1 = c$).

For the provider of dynamic contracts, the total profit can be expressed as

$$\mathrm{prof}_D(a,b) = \begin{cases} n\ b \int_0^{a/b} p(1-b)f(p)\,dp - c\ C(n,a,b)\ , & \text{if } a < b \\ n\ b \int_0^{1} p(1-b)f(p)\,dp - c\ C(n,a,b) & , \text{if } a \geq b\,, \end{cases}$$

where $C(n,a,b)$ is the size of the inventory, as a function of the number of customers acquired and the unit prices published by both providers.

Our first task is to define the capacity needed by a provider of dynamic contracts in order to fulfill the requirements of the n customers that reside in the market. Note that not all n customers will choose a dynamic contract, but we will address this issue later on. We use the theory of effective bandwidths in order to express the required total capacity C as a function of n and the expected bandwidth requirements by each user, given a QoS target. In particular, our QoS target defines the probability that the provider of dynamic contracts will not be able to honor all his contractual obligations, which can occur if at a particular slot the sum of the demands of his customers exceeds C, the size of his inventory (i.e. the demand will overflow the inventory). We call this the Inventory Overflow Probability (IOP).

We can express the QoS target with the following inequality:

$$P[\sum_i^n x_i \geq C] \leq e^{-\epsilon}\,, \tag{5}$$

where $e^{-\epsilon}$ denotes the IOP and x_i the bandwidth requested by customer i at a random slot.

From the theory of Chernoff bounds, we have that for a sequence of i.i.d. random variables $x_1, \ldots, x_n$ it holds:

$$P\left[\sum_i^n x_i \geq C\right] \leq E\left[e^{\theta(\sum x_i - C)}\right]\ , \forall\theta\,. \tag{6}$$

The right-hand side of the above inequality can be re-written as follows:

$$E\left[e^{\theta(\sum x_i - C)}\right] = e^{-n\left(\theta C/n - \log E\left[e^{\theta x_i}\right]\right)}\,. \tag{7}$$

Hence, we obtain that

$$P\left[\sum_i^n x_i \geq C\right] \leq e^{-n \sup_\theta \left[\theta C/n - \log E\left[e^{\theta x_i}\right]\right]}\,. \tag{8}$$

Combining (5) and (8), the sufficient condition for the QoS target to hold is

$$e^{-n \sup_\theta \left[\theta C/n - \log E\left[e^{\theta x_i}\right]\right]} \leq e^{-\epsilon} \Rightarrow \sup_\theta \left[\theta C - n \log E\left[e^{\theta x_i}\right] - \epsilon\right] \geq 0$$

$$\Rightarrow \sup_\theta \left[C - \frac{n}{\theta}\log E\left[e^{\theta x_i}\right] - \frac{\epsilon}{\theta}\right] \geq 0 \Rightarrow C - \inf_\theta \left[\frac{n}{\theta}\log E\left[e^{\theta x_i}\right] + \frac{\epsilon}{\theta}\right] \geq 0\,. \tag{9}$$

The above inequality defines the minimum capacity that the provider should have available in his access network in order to guarantee a QoS target of $IOP = e^{-\epsilon}$ in his contracts with the customers.

What remains to be defined is how we obtain x_i from our customer demand model in order to calculate the $E\left[e^{\theta x_i}\right]$ factor, from which we calculate the total capacity needed by solving the minimization problem in (9).

We have seen that if a customer's p_i is greater than a/b, then the customer will choose a static contract. This knowledge can be used in the calculation of the effective bandwidth. Thus, instead of considering all the customers in the market (n), we only examine those that prefer dynamic contracts (n'). It holds that $n' = na/b$, on average.

The distribution of bandwidth that a customer of dynamic contract requests is now more simple to express: with probability p_i he buys $1-b$ units of bandwidth, while with probability $1-p_i$ he does not buy nothing at all, with p_i being uniform in $[0, a/b]$. Hence, $E\left[e^{\theta x_i} | p_i = p \leq a/b\right] = pe^{\theta(1-b)} + (1-p)$.

Thus, we have that

$$\begin{aligned} E\left[e^{\theta x_i}\right] &= E\left[E\left[e^{\theta x_i} | p_i = p \leq a/b\right]\right] = \tfrac{b}{a}\int_0^{a/b}\left(pe^{\theta(1-b)} + (1-p)\right)dp \\ &= \left(\tfrac{a^2}{2b^2}e^{\theta(1-b)} + \tfrac{a}{b} - \tfrac{a^2}{2b^2}\right)\tfrac{b}{a} = \tfrac{a}{2b}e^{\theta(1-b)} + 1 - \tfrac{a}{2b}\,. \end{aligned}$$

Note that the factor b/a in the integral is actually the probability density function of p_i since now p_i is uniformly distributed in $[0, a/b]$. From the above, we get that the required capacity for the inventory that a provider of dynamic contracts needs to purchase, so as not to violate the IOP target is given by the following formula (note that n is now replaced by na/b):

$$C \geq \inf_{\theta}\left[\frac{\epsilon}{\theta} + \frac{a}{b}\frac{n}{\theta}\log\left(\frac{a}{2b}e^{\theta(1-b)} + 1 - \frac{a}{2b}\right)\right]\,. \tag{10}$$

5 Market analysis

In this section we present some of the most representative results using the previous model to analyze the behavior of the market, the equilibrium point, the profits of both providers and the market segmentation. Our purpose is to compare the case of provision affected by the number of customers in the market (especially when n is small) with the earlier case where we assumed constant costs, which corresponds to $n = \infty$.

More precisely, the analysis setup involves a price competition game between a provider offering static contracts and a provider offering dynamic contracts. The system parameters are the cost c for the provision of one unit of bandwidth and the number n of customers in the market. We assume that IOP is fixed and equal to 0.01. Each provider calculates a price as the best response (in terms of revenue maximization) at every possible price his competitor may post.

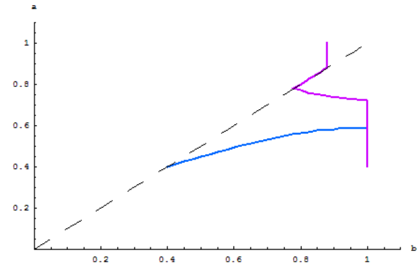

(a) $n = 10$

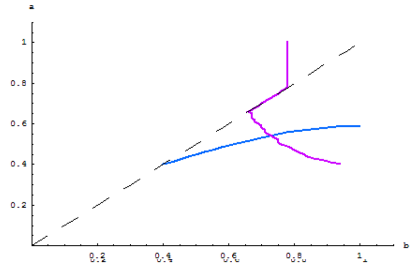

(b) $n = 55$

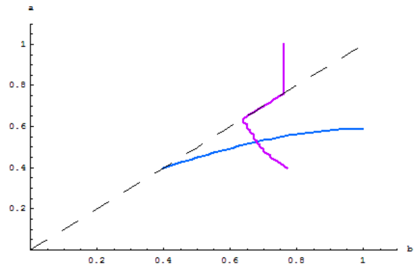

(c) $n = 100$

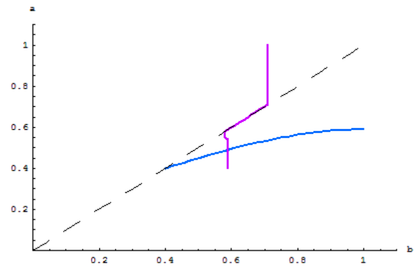

(d) $n = 3450$

Figure 4: The reaction curves $a(b)$ (blue) and $b(a)$ (violet) for $c = 0.4$. We obtain that for small n the provider of dynamic contracts does not participate in the market.

5.1 Reaction curves and equilibrium points

First of all, we examine the evolution of the reaction curves as the number of customers increases. It will be very interesting to see what happens for a small number of customers were multiplexing is not so efficient and effective bandwidths are large. Obviously, only the $b(a)$ reaction curve will change, since there is no change in the way the provider of static contracts decides his best response a to a published price b by the provider of dynamic contracts.

Figure 4 shows the reaction curves of the two providers. The blue curve depicts the reaction of provider that offers static contracts. The red curve depicts the reaction of the provider that offers dynamic contracts. The unit cost is set to be $c = 0.4$. The size of the market varies from very small ($n = 10$), to medium ($n = 55$ and $n = 100$) and very large ($n = 3450$).

When the size of the market is very small, the provider of dynamic contracts cannot take full advantage of multiplexing. Thus, the effective bandwidth per customer is very high and it is not beneficial for the dynamic provider to participate in such a market. This explains why, for $n = 10$, we get that $b = 1$ at the equilibrium point, which means that no customer will prefer a dynamic contract at such a high price. Recollect that the demand function for a single time slot is $1 - b$, leading to zero consumption when $b = 1$. As the size of the market increases, multiplexing becomes more efficient, provision costs decrease and the dynamic contracts become beneficial for some customers. Hence, at the equilibrium point it will hold that $b < 1$. Finally, when the market size becomes quite large, the respective reaction curves tend to become identical with the reaction curves of the initial model with fixed provision costs, where $n = \infty$ (compare with Fig. 2(a)). This happens due to the fact that the market is so large that the number of customers does not affect the provision cost any more, since their flows are ideally multiplexed and the effective bandwidth is equal to the average bandwidth a customer requests.

The new reaction curve $b(a)$ has three segments, similar to the reaction curve of Section 3. In the first segment, for very high values of a, both approaches provide the same, fixed value for b. Then, in the second segment, as a decreases, b also decreases and in fact it holds that $b = a$, for the same reason explained in Section 3, where we had that $n = \infty$. At the third segment, i.e. for lower values of a, b increases, since the provider of dynamic contracts has to cover his large provision costs and cannot continue decreasing b. This is also due to the fact that the probability for a customer to select a dynamic contract ($p < a/b$) is becoming lower. For a better understanding of how the profits under a dynamic contract change when b varies and a remains constant, consult Section 5.2.

As far as the existence of an equilibrium point is concerned, we have seen that as the size of the market changes, the equilibrium point gets closer to the equilibrium point calculated when having fixed unit costs (see Section 3). We have also explained what happens for small market sizes, where it is not beneficial for a provider who offers dynamic contracts to participate. We need to show whether the existence of an equilibrium depends on the value of c.

Figure 5 depicts how the equilibrium point moves as the provision unit cost increases, for a small market ($n = 20$). It is obvious that as the cost increases, all reaction curves are shifted upwards and the reaction curves of the dynamic provider are also shifted to the right. As a result, increasing unit costs result in higher prices and for very high costs, the provider of dynamic contracts has no benefit in participating in such a market. Simulations have also shown that for large-sized markets ($n > 2500$), a non-operational equilibrium point will be reached only for extremely high unit costs ($c > 0.9$). To illustrate this better, Fig. 6 shows the minimum market size n that makes it profitable for the dynamic provider to enter the market, as c increases. It is obvious that the minimum market size is growing exponentially, meaning that unit cost must acquire very high values so that even when the market is large market, it is not beneficial for the dynamic provider to be activated in such a market.

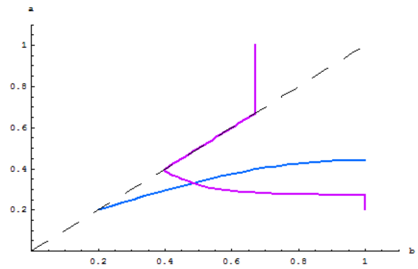

(a) $c = 0.2$

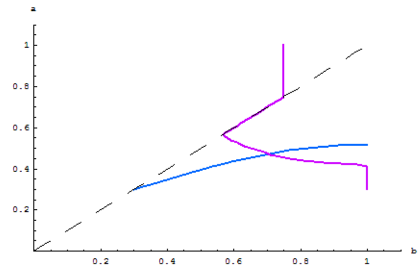

(b) $c = 0.3$

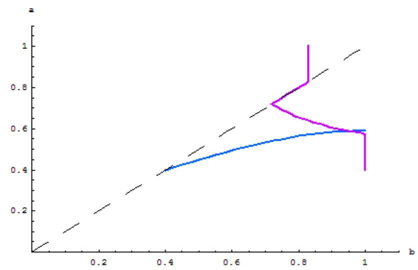

(c) $c = 0.4$

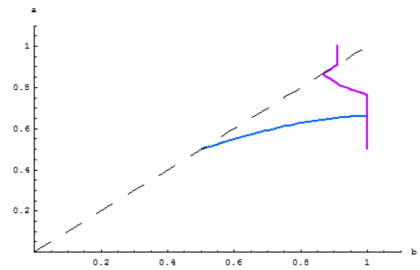

(d) $c = 0.5$

Figure 5: The reaction curves for the providers of static(blue) and dynamic(violet) providers, as c increases and $n = 20$. As costs increase the dynamic provider is discouraged from joining the market ($b \to 1$).

5.2 Provider's profits

An interesting property to examine is how the profits of the dynamic provider change as prices a and b change, in order to understand how the reaction curve $b(a)$ is formed. Our purpose is to investigate the properties of each segment of the $b(a)$ reaction curve so as to understand better its shape. As a reference, we will use Fig. 4(c), where $n = 100$ and $c = 0.4$.

Figure 7 shows how the profits change for fixed values of a, as b varies. The first obvious result is that while $b < a$, as b increases so do the profits of the dynamic provider. This forms the first segment of the profit curve, which is either linear or concave. When b becomes larger than a, we have the second segment of the profit curve, which is also linear or concave, depending again on the value of a. Hence, the profit maximizing point is either where the first and second segment meet, or strictly on the second segment.

For high values of a, the profit maximization point remains fixed (see Fig. 7(a) and 7(b)). This explains the first linear segment of the reaction curve $b(a)$. As seen in the figures, the maximizing point is where the two segments of the profit curve meet. As a continues to decrease (second segment of the reaction

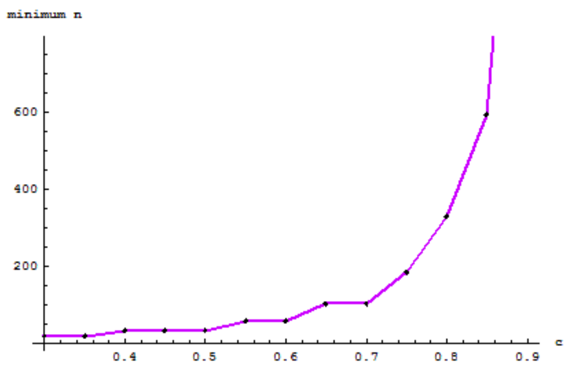

Figure 6: Minimum market size required for the provider of dynamic contracts to enter the market, as c increases. The exponential behavior of the minimum required market size is obvious.

curve), the maximizing b becomes equal to a and starts decreasing (see Fig. 7(c) and 7(d)). In these figures, we observe that the maximizing point is still at the meeting point of the two segments of the profit curve, but this point moves to the left, as a decreases, resulting in a smaller linear segment and a more concave second segment. There is a point where the maximizing value of b becomes the local maximum of the second segment. Thus, the optimal b starts increasing as a decreases (third segment of the reaction curve) (see Fig. 7(e) and 7(f)).

It is also interesting to see how the profits of the static and dynamic contracts are related, especially for small markets. It is expected that the larger the market, the higher the profits of the dynamic provider will be. In Fig. 8, we depict the providers' profits for small-sized market. As expected, for very small markets, the provider of static contracts obtains higher profits since the dynamic provider either does not participate in the market, or his unit price b is large enough in order to cover his provision costs.

So far, we have illustrated the reaction curve of the dynamic provider when in a price competition environment. We have also seen, how the profit maximization problem affects these reaction curves. In the next section, we will show what happens with the bandwidth provision considerations for the provider's bandwidth inventory.

5.3 Bandwidth allocation comparison

As mentioned earlier, the per customer bandwidth provision is expected to be higher than the average bandwidth per user and lower than the peak bandwidth

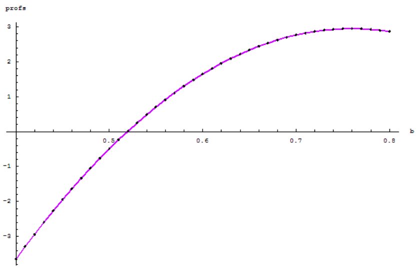

(a) $a = 0.8$

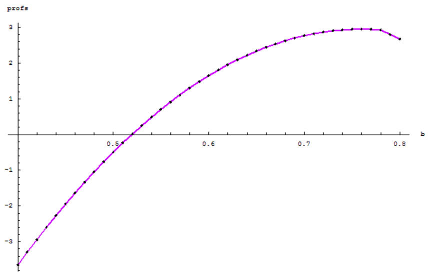

(b) $a = 0.78$

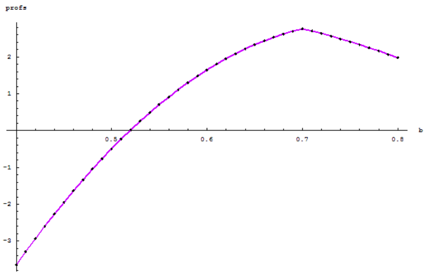

(c) $a = 0.7$

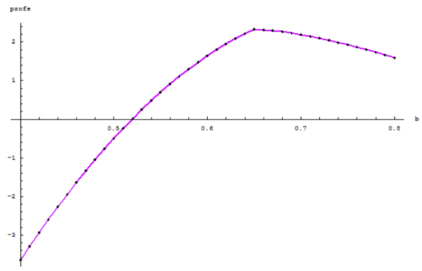

(d) $a = 0.65$

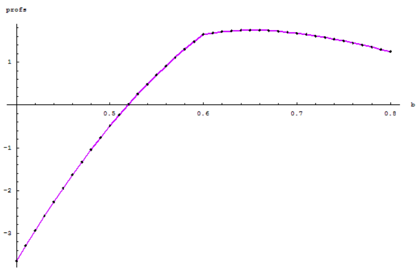

(e) $a = 0.6$

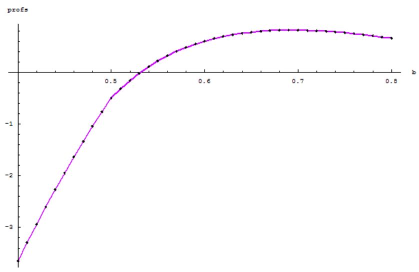

(f) $a = 0.5$

Figure 7: The profits curve for fixed values of a and varying b ($n = 100$). The maximum in each curve defines how the reaction curve $b(a)$ is formed and explains its three linear segments.

that a customer may require. In this section, we show that the numerical analysis of our model is validated by the theory of effective bandwidths.

First, we examine the amount of bandwidth that the provider of dynamic contracts needs to reserve in his inventory for each customer, when using peak, average and effective bandwidth reservation techniques. In Fig. 9, we keep a and b fixed and we only vary the size n of the market.

From the above diagram, it is obvious that the effective bandwidth is always

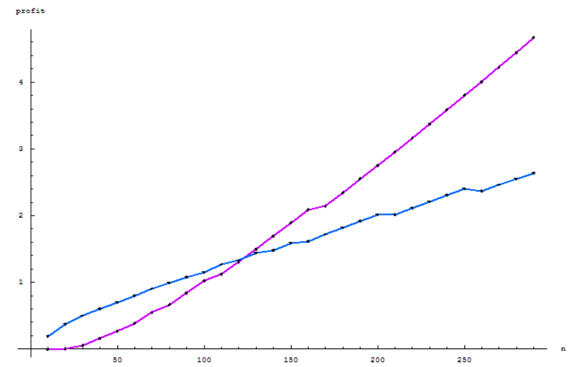

Figure 8: Total profits for the providers of static and dynamic contracts, as the market size increases. Violet line represents profits of the dynamic provider while blue line shows the profits of the static provider ($c = 0.4$). For small market sizes, static contracts are more profitable. The result is reversed for larger markets.

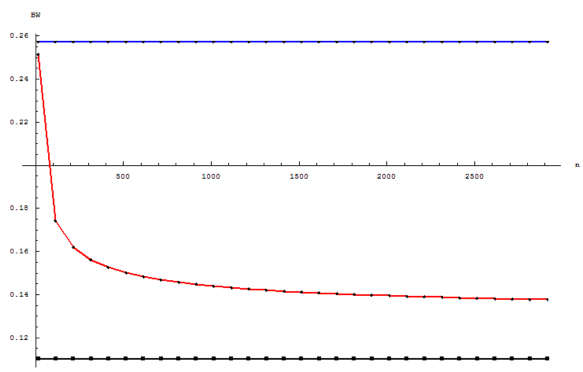

Figure 9: The per customer effective (red), average (black) and peak (blue) bandwidth needed to be reserved, for fixed prices and unit cost ($a = 0.6$, $b = 0.7$ and $c = 0.3$), as n increases. Observe that the effective bandwidth tends to become equal to the average, as n increases.

between the peak and average bandwidth, as expected. Furthermore, as the number of customers increases, the effective bandwidth converges to the average bandwidth. This result is in accordance with the theory of effective bandwidths. The variation of the unit cost affects only the actual bandwidth allocated and not the relative position of the curves.

5.4 Market Segmentation

One last issue, is how the market is segmented between the providers offering static and dynamic contracts. Recollect that, whether a customer chooses a dynamic or static contract, depends on his type, i.e. the value of p, and the unit prices of the market. So, if $p > a/b$ then the customer prefers a static contract. This preference is reversed when $p \leq a/b$. In Fig. 10, we plot the a/b value, for two unit costs ($c = 0.3$ and $c = 0.6$) as the market size changes.

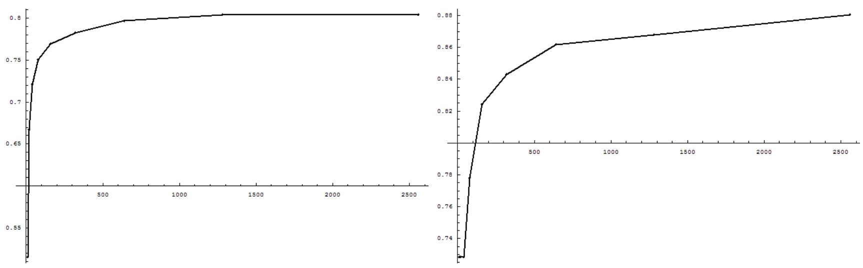

(a) $c = 0.3$ (b) $c = 0.6$

Figure 10: The market segmentation threshold (a/b) for fixed values of c, as n increases. The portion of customers preferring dynamic contracts is above half, even for smaller market sizes.

We observe that a provider of dynamic contracts obtains a large portion of the market. This portion increase as the market size increases and becomes almost fixed for large markets. Even in a situation with low provision costs, such a provider obtains more than 70% of the market share. And this happens even in relatively small markets (e.g. for $n = 100$).

6 Conclusions

In this paper, we have provided a model for analyzing an access bandwidth market, where two providers offering different types of contracts compete in a price game with the goal of maximizing their profits. One of the providers offers static contracts under which customers buy the necessary bandwidth at the beginning of a long-term period. The other provider offers dynamic contracts where each customer can express his demand at the beginning of smaller-scale periods (slots). Both providers seek to maximize their profits by posting prices at the beginning of time and keeping them fixed through the long-term period.

We have provided a model to capture the customers' varying demand for bandwidth of the customers, allowing them to have different levels of demand in each slot. We have calculated the reaction curves for the case of constant provision costs and we have shown the existence of an equilibrium point. Then considering the implications of the market size in the provision cost, we have tried to capture the effects of multiplexing in the cost structure of the dynamic provider. With the use of the effective bandwidths theory, we have extended our initial problem and have provided the new maximization problems. Through numerical analysis, we have plotted the reaction curves and we have investigated how they are affected by the market size. Furthermore, we have studied the conditions under which a provider of dynamic contracts can participate in the market, the profits he makes and the market segment he obtains.

As future work, we would like to investigate how different models for demand affect our results. More precisely, in our initial model of constant costs, we would like to see how the reaction curves are affected under different distributions for p. In the model of multiplexing-dependent provision cost for the dynamic provider, we would like to examine if and how we can include a penalty for the provider of dynamic contracts when not fulfilling the customer's bandwidth requirements, i.e. the probability defined by the IOP for the inventory overflow. Finally, the properties of *mixed contracts* must be studied in the case of a competitive environment. Mixed contracts are a combination of static and dynamic contracts, under which the customer can buy at a price a a static amount of bandwidth for the whole longterm period (static part), but has also the ability to buy extra bandwidth at a price b, at the beginning of each slot (dynamic part). In our previous work we have examined some of their properties in a monopoly setup.

Acknowledgments

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 has been supported by the EuroNGI Network of Excellence (EU-IST -507613) and has also a partial support of the IST project GridEcon (IST-2006-033634) funded by the EU.

References

[Binmore 1991] Binmore, K.: "Fun and Games - A Text on Game Theory," Lexington, MA: D. C. Heath and Co., Eds., 1991.

[Courcoubetis and Weber 2003] Courcoubetis, C., Weber, R.: "Pricing Communication Networks: Economics, Technology and Modeling," Chichester, Wiley, 2003.

[Courcoubetis et al. 2006] Courcoubetis, C., Soursos, S., Weber, R.R.: "Pricing Resources on Demand," *1st IEEE International Workshop on Bandwidth on Demand,* 2006, pp. 12-18.

[Davies et al. 2004] Davies, G., Hardt, M., Kelly, F.: "Come the revolution-network dimensioning, service costing and pricing in a packet switched environment," *Telecommunications Policy,* 2004, v.28, pp. 391-412.

[Foros and Hansen 2001] Foros, O., Hansen, B.: "Competition and compatibility among Internet Service Providers," Information Economics and Policy, 2001, vol. 13, issue 4, pp. 411-425

[Gibbens et al. 2000] Gibbens, R., Mason, R., Steinberg, R.:"Internet Service Classes under Competition," IEEE Journal on Selected Areas in Communications, 18(12), December 2000, 2490-2498

[Mason 2000] Mason, R.: "Simple Competitive Internet Pricing," *European Economic Review,* 44(4-6), 2000, pp. 1045-1056.

[Rabbat and Hamada 2006] Rabbat, R., Hamada, T.: "Revisiting Bandwidth-on-Demand Enablers and Challengers of a Bandwidth Market," *Network Operations and Management Symposium,* 2006, NOMS 2006, 10th IEEE/IFIP, pp. 1-12.