Economic Models of Sponsored Content in Wireless Networks with Uncertain Demand

Matthew Andrews
Alcatel-Lucent Bell Labs
Murray Hill, NJ

Ulaş Özen
Ozyegin University
Istanbul, Turkey

Martin I. Reiman
Alcatel-Lucent Bell Labs
Murray Hill, NJ

Qiong Wang
University of Illinois
Urbana-Champaign, IL

***Abstract*—The interaction of a content provider with end users on an infrastructure platform built and maintained by a service provider can be viewed as a two-sided market. Content sponsoring, i.e., charging the content provider instead of viewers for resources consumed in viewing the content, can benefit all parties involved. Without being charged directly or having it counted against their monthly data quotas, end users will view more content, allowing the content provider to generate more advertising revenue, extracted by the service provider to subsidize its investment and operation of the network infrastructure.**

However, realizing such gains requires a proper contractual relationship between the service provider and content provider. We consider the determination of this contract through a Stackelberg game. The service provider sets a pricing schedule for sponsoring and the content provider responds by deciding how much content to sponsor. We analyze the best strategies for the content provider and service provider in the event that the underlying demand for the content is uncertain.

Two separate settings are defined. In the first, end users can be charged for non-sponsored views on a per-byte basis. In the second we extend the model to the more common case in which end users purchase data quotas on a periodic basis. Our main conclusion is that a coordinating contract can be designed that maximizes total system profit. Moreover, the additional profit due to sponsoring can be split between the content provider and service provider in an arbitrary manner.

I. Introduction

One of the biggest trends in wireless communications over the past few years has been the explosion in demand for data services driven by the introduction of smartphones such as Apple's iPhone. Meeting this demand requires large investments in wireless capacity. However, for various socio-economic reasons, the price of basic service that a service provider can charge end users is fixed, often at a level that cannot generate enough revenue to pay for the cost of upgrades. Therefore, to stay on a sustainable path, the service provider needs to explore income from other sources, preferably in a mutually-beneficial manner.

Bundling and variable pricing are two traditional ways for revenue enhancement. In the first case, the service provider offers a single package that contains both high-revenue-per-bit services and low-revenue-per-bit services, and uses the gains from the former to cover the loss of the latter. In the second case (e.g. [5]), the service provider charges different fees for the same service provided at different times-of-day or different locations, effectively using pricing as a control device to fit demand within capacity while discriminating between users with different needs. The potential of either method ultimately rests upon pockets of end users, and hence is limited by the latter's budget.

In this paper, we investigate an approach whereby the service provider can tap into an alternative source of revenue, originating from sales of advertisements and channelled by the content provider in the form of sponsorship of viewing. The gain comes from removing inefficiency of the current arrangement under which the content provider derives profit from showing advertisements while end users pay the cost of viewing them. To maximize its profit, the provider naturally wants to increase the number of views of its content, but is unlikely to get help from end users who are wary about wasting their precious bandwidth quotas on embedded advertisements. The reluctance is strong not only because the use of bandwidth can be heavy (in the case when an advertisement comes in with rich video), but also because end users are uncertain and thus cannot control such use. Unlike voice service, which is charged by minutes, data service is sold in units of bytes and it is harder for the end users to interpret how many bytes will be consumed when they perform a web action. Typically when an end user clicks on a link the only indication they would get that the resulting web page is large is if it takes a long time to load.

We consider a solution that allows the content provider to "sponsor" its content so that it does not get charged to the end-users' monthly quotas. The arrangement removes end users' concern about paying an uncertain amount of bandwidth cost for carrying advertisements that are of little immediate value to them. As a consequence, more content will be accessed, not only because some of it is free but also because users are effectively given more quota. The content provider's profit increases as long as the cost of sponsoring stays below the new advertising revenue from increased viewing of its content. Moreover, the provider's image also improves as fewer users will think of it as an irresponsible party that pushes costly and worthless materials to them.

Content sponsoring also benefits the service provider by giving it the opportunity to charge content providers who, as the "Over-The-Top" companies, have a greater willingness to pay than end users. Income from this new source enables the service provider to recover the value of some of the mobile services that it is enabling and use that revenue to finance capacity expansion. To achieve this end, proper contracting

978-1-4673-1017-8/13/$31.00 ©2013 IEEE

between the service and the content providers is crucial: the service provider needs to design a contract that induces the content provider to participate in sponsoring and make decisions that improve the SP's profit. It turns out that the contract we propose not only does this but also maximizes the gain to the entire system, while at the same time transferring a significant share of the gain to the service provider itself.

Similar issues of contract design have been discussed in a different context in the marketing discipline under the banner of "channel coordination" [4], and have been widely addressed in the supply chain literature (e.g., see [1] for review). More recent work in [2] and [6] has started the consideration of coordinating contracts in the network setting. Our paper fits closely with this stream of thoughts. In particular, our contracting arrangement is equivalent to the stylized buyback contract discussed in both [4] and [1].

Previous work that specifically considers the option of content provider pricing includes [3] and [6]. (The first considers the problem in the network utility maximization (NUM) context while the second considers a situation where end users and content providers may not cooperate.) Our model differs from this prior work in that the underlying demand from the end users is uncertain. It is not simply a function of price and user/content provider utility. (In addition the NUM model assumes strictly-concave utility functions as a consequence of demand being elastic with respect to price, and hence does not capture a natural situation where the content provider is paid a fixed price per view by advertisers.) Lastly, our model extends beyond simple per-byte pricing and attempts to capture the notion of end-user quota dynamics.

The fact that we treat underlying demand as a random variable has two effects on our analysis. First, the problem faced by the content provider is similar to the "Newsvendor" problem that is common in supply chain analysis. (In a Newsvendor problem a retailer must purchase inventory to cover uncertain demand over a fixed time period. When the time period is over the inventory only has a salvage value that is well below the purchasing price.) In our setting the Newsvendor problem arises because the content provider must decide how much content to sponsor in a time period without knowing what the end user demand is. Second, the uncertain demand coupled with a reservation fee paid in advance will allow the service provider to control how much content is sponsored by the content provider, even when the latter's revenue grows in proportion to the number of views of its content.

We model the interaction between the service and the content providers as a Stackelberg game in which the service provider offers a contract parameterized by two fixed fees: a reservation fee proportional to the maximum number of views to be sponsored, and a usage fee for each sponsored view that actually takes place. By accepting this contract, the content provider determines the maximum number of sponsored views and pays the corresponding reservation fee in advance, and assumes the payment of the usage fee for each view of its content by end users, up to the aforementioned maximum.

We divide our discussions into two parts that reflect different ways of modeling end user payments.

- In Section II we present a simple model in which service-provider congestion costs, end-user bandwidth costs, and the price that the content provider must pay for sponsoring content are all determined on a per-byte basis. A key feature of this model is that the underlying demand from the end users is a random variable but the service provider would like to control the amount of bandwidth it has to provider.
  We focus on the relationship between the service provider and a single content provider. We show that the aforementioned two-fee contract is incentive compatible: by charging a proper reservation fee, the content provider will be induced to choose the maximum number of sponsored views to optimize the total expected profit of both parties. It is also in the service provider's best interest to charge such a fee to bring about this outcome, since it can then use the per-use fee to transfer the profit to itself. In Section II-E we present a numerical example to demonstrate how the optimization might work in practice.
- In Section III we indicate how the results can be adapted for the case of end-user quotas. In most current wireless data plans the end users pay a certain fee for a fixed quota of data. They do not pay on a byte-by-byte basis. This has a significant effect on the model since it is now much less explicit how much end-user revenue the service provider is giving up when content is sponsored.

A. The Models

In this paper we consider a contractual relationship between a single service provider (SP) and a single content provider (CP) in offering sponsored views of content. The situation is formally modeled as a Stackelberg game in which the SP is the leader who sets price parameters of the contract and the CP is the follower who responds by determining the maximum number of views it is willing to sponsor within a fixed period, e.g., monthly. The purpose of sponsoring is to raise the advertising revenue by the increase of the number of views of the said content. To model this effect, we assume end users (EU) will always access the said content if it is sponsored and with a smaller probability if it is not.

This problem can be naturally extended to the case of multiple CPs competing for the attentions of the EUs by sponsoring their own content. That situation leads to interesting competitive dynamics between the content providers and we leave its analysis for future work.

We define two basic models. In the first the end users pay for bandwidth on a per-byte "pay-as-you-go" basis. In the second model we aim to capture the more common situation in which end users pay for bandwidth via monthly quotas.

a) Model of Sponsored Content with Per-Byte End User Costs: The EUs generate N (a random variable) potential views in a period for content items that for ease of exposition are all assumed to have the same size θ. (It would not be difficult to extend to a situation where θ is the mean size of heterogeneous content.) Let F be the cumulative distribution function of N and let $\bar{F} = 1 - F$. If the content is sponsored it is viewed with probability 1. If content is not sponsored it is viewed with probability $q < 1$. (The parameter q there captures the strategic behavior of the EUs.) Let $bin(m, q)$ denote a binomial random variable with $m \in \mathbb{N}^0$ trials and success probability q. Then $E[bin(m, q)] = mq$.

The SP charges only the CP for sponsored content and the EUs for non-sponsored content, all on a per-byte basis. As we assume constant size for content, we denote EUs' payment per view for non-sponsored content by r. The CP's decision is denoted by B, defined as the maximum number of views the CP is willing to sponsor. The actual number of sponsored views is therefore $\min\{N, B\}$ and the total number of views is $\min\{N, B\} + bin([N - B]^+, q)$. As mentioned earlier, the payment from the CP to the SP is structured as an ex ante reservation fee and an ex post usage fee. We define c to be the reservation fee per view and b to be the usage fee per view. Hence the total revenue that the SP collects from the CP in a given period is $cB + b\min\{N, B\}$.

It remains to define the advertising revenue earned by CP and the bandwidth cost incurred by SP. We assume that CP earns revenue a for each view and hence its total revenue is $a(\min\{N, B\} + bin([N - B]^+, q))$. (We assume that all parameters are known by the SP and leave the interesting case where a and q are private information of the CP for future work.) The cost for the SP is dependent on the total congestion on its network, which is given by a non-decreasing function $C(\cdot)$. We let $\underline{B}$ be the total load on the network excluding the EU's views of the CP's content. The total load without sponsoring is

$$B_0 = \underline{B} + \theta bin(N, q)$$

and the total load with a sponsoring level of B is

$$\underline{B} + \theta min[N, B] + \theta bin([N - B]^+, q).$$

So the expected congestion cost paid by the SP is

$$E[C(\underline{B} + \theta min[N, B] + \theta bin([N - B]^+, q))]. \tag{1}$$

Our goal is to determine the maximum number of views that the CP should sponsor to maximize the total profit of both CP and SP. We also study the fees charged by SP that can induce this outcome.

b) Model of Sponsored Content with EU Quotas: We now describe a more refined model in which EUs do not pay for bandwidth on a per-byte basis. Each EU instead pays periodically for a base data quota and has the ability to buy additional quota in case the base quota is exhausted in a period. More formally, we assume there is a homogeneous population of K EUs, all of whom are served by a single service provider (SP) who periodically charges a fixed subscription fee. At the beginning of each billing cycle, every EU gets a bandwidth quota that she can use anytime within the period. The starting point of the first cycle of EUs is uniformly distributed over a period length. When an EU has exhausted her quota before the end of a period, she can wait until she gets new quota at the beginning of the next period, or refill her quota immediately by paying an additional amount d. The choice between waiting and refilling is assumed to be independent of the number of times that the EU has refilled before.

An EU's opportunity to access content within a unit time period is a Poisson distributed random variable with mean λ. The likelihood of an EU taking the opportunity to view the content depends on the amount of the quota she has for the remainder of the period (this is a strong assumption since it does not take into account that an EU may use her quota more aggressively when it is about to expire). We model EU's decision by a discrete-state Markov chain. States are indexed by $i = 0, ..., S$, where EUs in the states of smaller index have more available quota left. EUs in state S have exhausted their quotas and are waiting for the next period to arrive. An EU in state i views unsponsored content with probability q_i ($i = 0, ..., S$), where $q_S = 0$. We remark that by using the Markov model the periods will not have an equal length. However, we focus on this model as an approximation to a regular billing cycle due to its tractability.

II. Analyzing Sponsored Content when EUs Pay Per Byte

A. Content provider's problem

We start with the simplest situation in which the reservation fee $c = 0$. Recall that B denotes the maximum number of sponsored views, $a > 0$ denotes the (advertising) revenue to the CP of each view, and $b \geq 0$ denotes the usage-fee per sponsored view paid by the CP to the SP. The revenue received by the CP is $E[a(\min(N, B) + bin([N - B]^+, q)]$ and the cost paid by the CP to the SP is $E[b\min(N, B)]$. The net revenue to the CP is

$$\begin{aligned} & E[(a - b)\min(N, B) + abin([N - B]^+, q)] \\ = \; & aqE[N] + (a\bar{q} - b)E[\min(N, B)] \end{aligned} \tag{2}$$

where $\bar{q} = 1 - q$, and we have used $[x - y]^+ = x - min(x, y)$. Let $B^*(b)$ denote the maximizing value of B for given b. Then $B^*(b) = \infty$ if $a\bar{q} > b$ and $B^*(b) = 0$ if $a\bar{q} < b$. If $a\bar{q} = b$, the CP's net revenue function becomes a constant and hence the CP is indifferent between any choice of $B^*(b) \in [0, \infty)$. In other words, if the CP does not need to pay a reservation fee in advance for sponsoring but price b is paid for each view of sponsored content, then the CP's optimal choice of $B^*(b)$ is either zero or infinity with a transition point where the CP is indifferent between sponsoring any content or not.

Now, consider the case with a per-unit reservation fee $c > 0$. Then CP's revenue function becomes

$$aqE[N] + (a\bar{q} - b)E[\min(N, B)] - cB.$$

This is a standard Newsvendor model. If $c \geq a\bar{q} - b$ then $B^* = 0$, i.e., the CP will not sponsor any content viewing if the combined reservation and usage fees exceed the additional revenue from advertisement. If N has a continuous distribution function (F has no jumps) then

$$B^* = \bar{F}^{-1}(\frac{c}{a\bar{q} - b}). \tag{3}$$

In our setting N is a discrete random variable, so F has jumps, and there may not be a B^* such that (3) holds exactly. On the other hand, N is likely to be an extremely large integer (in our numerical example $E[N]$ is on the order of 10^7), so (3) will hold *almost* exactly. In particular, since in general B^* is defined by

$$\bar{F}(B^*-1) > \frac{c}{a\bar{q}-b}, \quad \bar{F}(B^*) \le \frac{c}{a\bar{q}-b},$$

the error we make in assuming that (3) holds is miniscule and will henceforth be ignored.

B. Service provider's problem

First, consider the case with no contract cost (i.e. $c = 0$) and no revenue from end users (i.e. $r = 0$). The SP's revenue from the CP is $bE[\min(N, B)]$. The SP pays a congestion cost, given by a function C. Recall that $\underline{B}$ is the "baseline" congestion without the EU, and the congestion cost is given by (1). We remark that congestion cost may not be a convex function of B even if C is linear since $\min[N, B]$ is concave.

The SP wants to choose b to maximize

$$\begin{aligned} & bE[\min(N, B^*(b))] \\ - \; & E[C(\underline{B} + \theta min[N, B^*(b)] + \theta bin([N - B^*(b)]^+, q))]. \end{aligned}$$

Recall that with $b < a\bar{q}$, $B^*(b) = \infty$, and with $b > a\bar{q}$, $B^*(b) = 0$. Thus the SP wants to choose either b as large as possible subject to $b < a\bar{q}$, in which case the SP's profit is

$$\begin{aligned} & \lim_{b\uparrow a\bar{q}} bE[\min(N, B^*(b))] \\ - \; & E[C(\underline{B} + \theta min[N, B^*(b)] + \theta bin([N - B^*(b)]^+, q))] \\ = \; & a\bar{q}E[N] - E[C(\underline{B} + \theta N)], \end{aligned}$$

or $b > a\bar{q}$, in which case the SP's profit is $-C(\underline{B} + \theta bin(N, q))$. The SP will choose the alternative yielding the higher profit.

Now consider the case with fixed contract cost $c > 0$. The SP's revenue from the CP is $bE[\min(N, B)] + cB$. So the SP wants to choose c and b to maximize

$$\begin{aligned} & bE[\min(N, B^*(b, c))] + cB^*(b, c) \\ - \; & E[C(\underline{B} + \theta min[N, B^*(b, c)] \\ + \; & \theta bin([N - B^*(b, c)]^+, q))]. \end{aligned}$$

Given the relationship $\bar{F}(B^*(b, c)) = \frac{c}{a\bar{q}-b}$ from (3), the SP's problem is equivalent to choosing $B \ge 0$ and $b \in [0, a\bar{q})$ to maximize

$$\begin{aligned} & bE[\min(N, B)] + \bar{F}(B)(a\bar{q} - b)B \\ - \; & E[C(\underline{B} + \theta min[N, B] + \theta bin([N - B]^+, q))]. \end{aligned}$$

We now consider the optimal b for a given B. Looking at the first order derivative of the profit function

$$E[\min(N, B)] - \bar{F}(B)B = \int_0^B \bar{F}(x)dx - \bar{F}(B)B \ge 0,$$

we conclude that SP wants to set b as high as possible such that $b < a\bar{q}$. Thus, with B fixed the SP's profit is

$$\begin{aligned} & \lim_{b\uparrow a\bar{q}} bE[\min(N, B)] + \bar{F}(B)(a\bar{q} - b)B \\ - \; & E[C(\underline{B} + \theta min[N, B] + \theta bin([N - B]^+, q))] \\ = \; & a\bar{q}E[\min(N, B)] \\ - \; & E[C(\underline{B} + \theta min[N, B] + \theta bin([N - B]^+, q))], \end{aligned}$$

so that the optimal profit is attained using B^* given by

$$\begin{aligned} B^* \; = \; & \arg\max_B \{a\bar{q}E[\min(N, B)] \\ & -E[C(\underline{B} + \theta min[N, B] + \theta bin([N - B]^+, q))]\}. \end{aligned}$$

Of course, in practice the limit cannot be attained: the SP needs to keep $c > 0$ to induce the CP to choose the SP's desired B^*. Thus there is some small ε such that $b = a\bar{q} - \varepsilon$ and $c = \bar{F}(B^*)\varepsilon$, i.e., while a positive reservation fee is necessary to induce optimal B, the SP is better off to keep it as low as possible and derive all its profit by setting b as high as possible.

We also remark that due to equation (3), when finding B^* we must optimize over the support of N. In reality we would typically wish to restrict the optimization further to between (say) $F(0.02)$ and $F(0.98)$ since otherwise the system would be overly senstive to the exact values of b and c.

Revenue from EUs: Suppose that SP earns revenue from the EUs, i.e., $rE[bin((N - B)^+, q)]$ where r is the revenue rate. So the SP wants to choose c and b to maximize

$$\begin{aligned} & bE[\min(N, B^*(b, c))] + cB^*(b, c) \\ - \; & E[C(\underline{B} + \theta min[N, B^*(b, c)] + \theta bin([N - B^*(b, c)]^+, q))] \\ + \; & rE[bin((N - B^*(b, c))^+, q)]. \end{aligned}$$

If $r \ge \frac{a\bar{q}}{q}$, it is not beneficial for the SP to offer sponsored content option to CP since the additional revenue from CP is not high enough to compensate for the loss in revenues from the end users. This is the case if ad revenues are low (i.e., for low a values) and/or the content is popular (i.e., for high q values).

Now, consider the case with $r < \frac{a\bar{q}}{q}$. Similarly as above, the SP's problem is equivalent to choosing B and b to maximize

$$\begin{aligned} & bE[\min(N, B)] + \bar{F}(B)(a\bar{q} - b)B \\ - \; & E[C(\underline{B} + \theta min[N, B] + \theta bin([N - B]^+, q))] \\ + \; & rE[bin((N - B)^+, q)]. \end{aligned}$$

Checking the derivative with respect to b, we conclude that SP wants to set b as high as possible such that $b < a\bar{q}$ as above. We can again optimize over B, the only difference being the additional term $rE[bin((N - B)^+, q)]$.

C. A Pareto analysis of the two-parameter contract

The system performance, i.e. the aggregate profit achieved by the SP and the CP, is given by

$$\begin{aligned} \pi^S(B) \; = \; & E[a\min(N, B) + (a + r)bin([N - B]^+, q)] \\ & -E[C(\underline{B} + \theta min[N, B] + \theta bin([N - B]^+, q))], \end{aligned}$$

and the SP takes the following share

$$\begin{aligned} & bE[\min(N,B)] + \bar{F}(B)(a\bar{q}-b)B \\ - \quad & E[C(\underline{B} + \theta min[N,B] + \theta bin([N-B]^+,q))] \\ + \quad & rE[bin((N-B)^+,q)]. \end{aligned}$$

Let B^S denote the system optimal number of sponsored views, i.e., $B^S = \arg\max_B\{\pi^S(B)\}$. The increase of the total expected profit from sponsoring is

$$\pi^S(B^S) - (r+a)qE[N] - C(\underline{B} + \theta bin(N,q)),$$

and from the system's perspective, sponsoring only makes sense if

$$\pi^S(B^S) > (r+a)qE[N] - C(\underline{B} + \theta bin(N,q)).$$

From the earlier discussion, we know that the inequality is not satisfied if $r \geq \frac{a\bar{q}}{q}$. However, we remark that even if $r \leq \frac{a\bar{q}}{q}$, sponsored content might not generate sufficient advertising revenue to offset the combined effect of losing EU revenue and increasing congestion cost. Such situations will be identified by the optimization of $\pi^S(B)$ when the optimal solution $B^S = 0$.

Assume that $\pi_S(B^S) \geq (r+a)qE[N] - C(\underline{B}+\theta bin(N,q))$. Since the system profit is maximized at B^S and the SP's share is increasing with b for any given B (and hence CP's share is decreasing with b), a contract (b,c) is Pareto efficient if and only if $B^*(b,c) = B^S$. Therefore, Pareto efficient contracts can be characterized by single parameter $b < a\bar{q}$ where $c = \bar{F}(B^S)(a\bar{q} - b)$. Under the set of Pareto efficient contracts, any allocation of additional system profit can be possible. The SP's share of profit will have a range of $[rqE[N] - C(\underline{B} + \theta bin(N,q)), \pi^S(B^S) - aqE[N])$ whereas CP's profit will have a range of $[aqE[N], \pi^S(B^S) - rqE[N] + C(\underline{B}+\theta bin(N,q))]$. We remark that for CP to achieve profit of $\pi^S(B^S) - rqE[N] + C(\underline{B} + \theta bin(N,q))$, CP needs to set b such that

$$\begin{aligned} & rqE[N] - C(\underline{B} + \theta bin(N,q)) \\ = \quad & bE[\min(N,B^S)] + \bar{F}(B^S)(a\bar{q}-b)B^S \\ - \quad & E[C(\underline{B} + \theta \min[N,B^S] + \theta bin((N-B^S)^+,q))] \\ + \quad & rE[bin((N-B^S)^+,q)]. \end{aligned}$$

When r is small, b that satisfies this equality might be negative.

D. Summary of the analysis with a contract price c and additional revenue from end users

The findings from the above analysis can be summarized as follows.

- The system performance, i.e. the total expected profit for both the SP and the CP, is given by

$$\begin{aligned} \pi^S(B) \quad = \quad & E[a\min(N,B) + (a+r)bin([N-B]^+,q)] \\ - \quad & E[C(\underline{B} + \theta min[N,B] + \theta bin([N-B]^+,q))]. \end{aligned}$$

  Optimizing the above determines whether sponsoring should take place.
- Although, as noted above, the system profit function $\pi^S(B)$ is not necessarily a concave function, finding the system optimal B is not a hard problem given that $\pi^S(B)$ is defined by one variable.
- Charging a single usage fee b is not enough to enforce a sponsored view to the CP. SP needs to charge a reservation fee c to induce the optimal limit on the sponsored views. By keeping c as low as possible and b as high as possible, the SP transfers all the expected gains from sponsoring to itself.
- Under any coordinating contract, the system profit is maximized (i.e., achieves Pareto optimum of the CP and the SP profits) Moreover, any allocation of additional profits is possible. Such a contract is definitely a win-win-win contract for CP, SP and end users.

E. Numerical Example

We now present a numerical example to illustrate the above concepts. Consider the case of a large CP for which N has a truncated normal distribution with mean $\bar{N} = 5 \times 10^7$ views per month. (The distribution is truncated to two standard deviations on each side.) The size of the content θ is 7.416Mbits and the CP receives \$0.0125 profit for each view (before paying any sponsoring charge to the SP). We assume end users pay at a rate \$10 per GB for non-sponsored content and so this translates to a cost per view of $r = 10\theta/(8 \times 10^9)$. For the SP congestion cost we set the baseline congestion $\underline{B} = 0$ and use a piecewise linear function given by,

$$C(x) = \begin{cases} 3rx/5\theta & \text{if } x \leq 19\bar{N}/20 \\ 10r(x - 893\bar{N}/1000)/\theta & \text{otherwise.} \end{cases}$$

(This stylized cost function reflects, in a simple manner, the additional costs, such as lost customer good will, of exceeding the nominal system capacity.) We set $q = 0.2$, i.e. an end user is five times as likely to view the content when it is free to them than they are when they have to pay for the bandwidth. In Figures 1 and 2 we show system profit as a function of B when the standard deviation of the underlying normal distribution is $2\bar{N}/5$ and $\bar{N}/10$ respectively. We can see that as the uncertainty in N increases (i.e. the standard deviation increases), the optimal amount of content to sponsor decreases since there is more likelihood that the realization of N will correspond to the steep part of the SP congestion cost curve. We can also see that although the system profit is not concave, it is simple to identify the optimal value of B. In

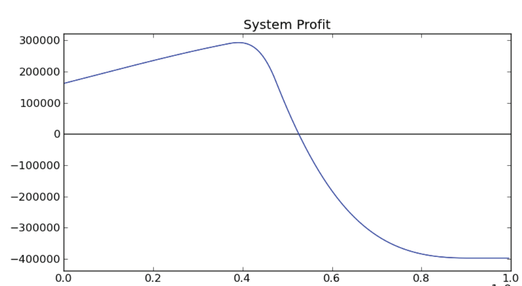

Fig. 1. System profit when standard deviation is $2\bar{N}/5$.

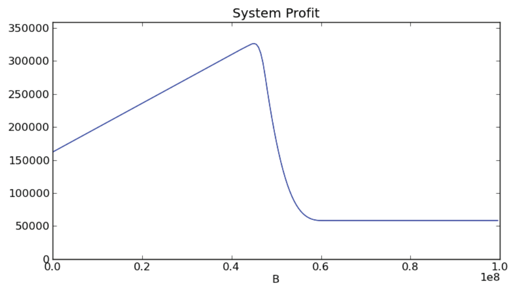

Fig. 2. System profit when standard deviation is $\bar{N}/10$.

Figure 3 we fix B to its optimal value (in this case 4.6×10^7 views) for the case that the standard deviation is $\bar{N}/10$. We then plot both SP profit and CP profit as a function of b. (Recall that c is then determined from b and B via Equation (3)). As b increases the excess system profit generated from the sponsored content is transferred from CP to SP.

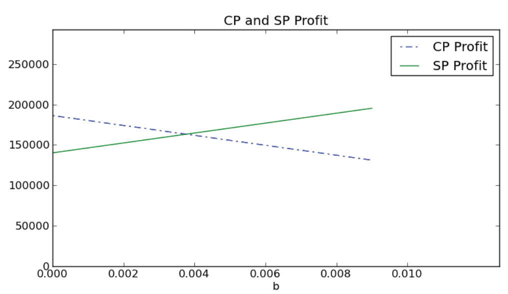

Fig. 3. SP and CP profit as a function of b when standard deviation is $\bar{N}/10$ and B is optimized for system profit.

III. Analyzing Sponsored Content in the Case of EU Quotas

Recall that we model quota usage via a discrete-state Markov Chain indexed by $\{0, \ldots, S\}$. Let P_i denote the steady-state probability that a user is in state i. Let q_i $(i = 0, \ldots S-1)$ denote the probability that a user in state i will view an unsponsored content. To properly represent the number of views we need to consider a multinomial distribution with $2(S+1)$ outcomes, corresponding to the $S+1$ states of the EU Markov chain along with viewed/not viewed. We index the outcomes $i = 0, \ldots, 2S+1$. Outcomes $i,\ i = 0, \ldots, S$ correspond to a view with the EU in state i. The probability associated with this outcome is $P_i q_i$. Outcomes $S+1+i,\ i = 0, \ldots, S$, correspond to no view with the EU in state i, and have associated probabilities $P_i(1-q_i)$. The total number of potential views is given by the random variable Λ^K, which has a Poisson distribution with mean λK. With the outcome probabilities understood, we let $N_i(M)$ denote the number of outcome i with M total trials. We assume that the EUs evolve independently and that viewing decisions are independent. Let $Q = \sum_{i=0}^{S-1} q_i P_i$, and let $\bar{Q} = 1 - Q$. Note that $Q < 1$.

Let $\underline{D}$ be the base revenue that SP receives from the EUs for their regular monthly quotas and let τ be the rate (which can be derived from the Markov Chain transition probabilities) at which EUs refill their quota "early". Hence the SP revenue from the EUs is $\underline{D} + d\tau$.

For reasons of space we do not provide the full details of how the results are affected by end-user quotas. However, for the case in which the Markov Chain transition probabilities do not change when CP's content is sponsored (i.e. the users simply switch their viewing from another content provider) we can obtain similar conclusions as before with $\bar{Q}$ playing the role of $\bar{q}$. In particular, the CP decision leads to a relationship of the form,

$$B^*(b,c) = \begin{cases} \bar{F}^{-1}\left(\frac{c}{a\bar{Q}-b}\right) & \text{if } c < a\bar{Q} - b \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

and so the SP choice of b and c is in fact equivalent to a choice of b and B so long as $c > 0$. (Once again therefore we need $c > 0$ in order for the SP to be able to control the system.)

Similar to the per-byte cases the system profit for a fixed value of B is,

$$\begin{aligned} & a\bar{Q}E[\min(\Lambda^K, B)] + (\underline{D} + d\tau) \\ - \quad & E[C(\underline{B} + \theta \min[\Lambda^K, B]) + \theta \sum_{i=0}^{S} N_i([\Lambda^K - B]^+)]. \end{aligned}$$

Hence as in the previous model we can optimize system profit via a univariate optimization over B. Once the optimal value of the B has been obtained the split between the SP and CP can be controlled by an appropriate choice of b.

For the case in which the transition probabilities do depend on B, the situation becomes more complex. In particular, a larger value of B is likely to reduce the rate at which quota is consumed. Hence in the system optimization we need to trade off increased advertising revenue from additional views with the reduction in end-user revenue that would be obtained from refilling quotas and the increased SP congestion costs. The formula for system profit is similar to the above. However, the optimization becomes more involved since $\bar{Q}$ and τ are now functions of B. For reasons of space we defer an exact characterization of the various tradeoffs to the full version of the paper.

References

[1] G.P. Cachon. Supply chain coordination with contracts. *Handbooks in operations research and management science*, 11:229–340, 2003.

[2] L. He and J. Walrand. Pricing and revenue sharing strategies for internet service providers. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, volume 1, pages 205–216. IEEE, 2005.

[3] Hande P., Chiang M., Calderbank A., and Rangan S. Network pricing and rate allocation with content provider participation. In *INFOCOM*, pages 990–998, 2009.

[4] B.A. Pasternack. Optimal pricing and return policies for perishable commodities. *Marketing science*, 4(2):166–176, 1985.

[5] Ha S., Sen S., Joe-Wong C., Im Y., and Chiang M. Tube: time-dependent pricing for mobile data. In *SIGCOMM*, pages 247–258, 2012.

[6] Y. Wu, H. Kim, P.H. Hande, M. Chiang, and D.H.K. Tsang. Revenue sharing among isps in two-sided markets. In *INFOCOM, 2011 Proceedings IEEE*, pages 596–600. IEEE, 2011.