On the Economics of Multicasting

Yuval Shavitt
Bell Labs 4G-627
101 Crawfords Corner Rd.,
Holmdel, NJ 07733,
shavitt@ieee.org

Peter Winkler
Bell Labs 2C-365
700 Mountain Ave.,
Murray Hill, NJ 07974,
pw@bell-labs.com

Avishai Wool
Bell Labs 2A-338
700 Mountain Ave.,
Murray Hill, NJ 07974,
yash@acm.org

Abstract— **A supplier of multicast information services will often be faced with the following problem: Broadcasting to the whole customer base (including non-paying customers) is cheaper than multicasting only to the paying customers. However, broadcasting discourages potential customers from paying. The result is an economic game in which the supplier tries to maximize profit in the face of rational, but not omniscient, behavior by customers.**

In this work we build a model for such environments, which we believe is both reasonably realistic and amenable to mathematical analysis. The supplier's basic strategy is to broadcast every service for which the fraction of subscribed customers exceeds some threshold. We then model the customers' behavior, taking into account the possibility of customers receiving services for free. From this model, coupled with some mild assumptions on the supplier's cost structure, we can find the optimal setting of the supplier's broadcast threshold. The solution necessarily depends on choosing functions which describe the customers' utility for the offered services; we study in detail several such choices.

In all the examples we studied, our model predicts that the supplier's profits will be maximized if the supplier's broadcast threshold is set below 100%: The loss in revenue due to customers subscribing to fewer services is offset by the cost savings made possible by broadcasting the most popular services to all customers. We found our model to be fairly robust with respect to parameter choices. As such, we believe it can be of value to a supplier in devising a multicast/broadcast strategy, and that broadcasting when subscriptions are sufficiently high is likely to be the approach of choice in maximizing profits.

I. Introduction

A. Background

In services such as cable or satellite TV, the subscription model has traditionally been periodic; e.g., customers pay a monthly fee and receive the channels they subscribe to over the next month. As service demand becomes more diverse, more applications appear where the subscription model is no longer periodic and applies to shorter time frames. For instance, pay-per-view services allow customers to pay for a single movie or sport event. In the Internet, it is likely that similar multimedia and data services will be available, and customers will expect to be able to pay only for the service they want, in small quanta of time (video clips, stock quotes, etc.). In order to support such a customized subscription model, the information supplier needs to use some type of multicast capability.

The supplier has several options of how to implement her multicast capability, which depend on the underlying infrastructure she is using. If the infrastructure is a broadcast channel, (e.g., transmitting via satellite or cable), then multicast is typically implemented by encrypting the content; the required decryption keys are given only to customers who have payed their subscription fees [MQ95]. If the infrastructure is the Internet, then multicast can be implemented by multiple individual unicasts; or by using a newly-created, dedicated, IP-multicast group (cf. [DC85], [AFM92]), whose members are only the paying customers; or by reverting to encryption, coupled with appropriate key management (cf. [WGL98], [WHA98], [BMS99]).

All these methods of implementing multicast share the following property: Unless the group of subscribers is very small, broadcasting to the *whole* customer base is typically cheaper than incurring the operation costs associated with multicasting only to the paying subscribers of a particular service. This observation becomes more acute when the customers are allowed to personalize their subscription to a high degree of granularity. Therefore, from a cost perspective, when the number of paying subscribers is high, the supplier has an economic incentive to broadcast to everyone, and to tolerate some free-riders.

However, allowing free-riders is risky. When the customers realize that they occasionally receive services they did not pay for, their willingness to pay will drop, which may lead fewer subscribers and to a loss in revenue for the supplier. At the extreme, if the supplier always uses broadcast, rational customers will never pay.

Thus we arrive at an economic question: What is the strategy that maximizes the supplier's profit? Would she gain or lose by using broadcast at all, and under what circumstances? In order to answer these questions, we need to model a rational customer's behavior, and to analyze what is the strategy that would maximize his utility. Naturally, the strategies of the supplier and customers depend on each other, and form an economic game.

B. Related Work

The interest in Internet economics has been growing since the early 1990s. Particular attention has been given to topics such as pricing, and economic relations among ISPs (see [MB97] and the references therein).

As for multicasting, the aspect of fair cost sharing among the receivers has been analyzed in [HSE97], using the theory of cooperative games. Their model fits a scenario in which the multicast group is long-lived, and the shared multicast tree is owned by the receivers (or by a non-profit entity). In contrast, in our model the multicast groups are transient and short-lived. Furthermore, it is the profit-seeking supplier who bears the operating cost of the multicast infrastructure.

Game theory has also been used in less traditional and more technical problem domains, such as routing [LO97], [KVA98], [KO99], flow control [KL95], [AB98], and congestion control [She95]. In these works the decentralized control of networks is modeled as a game between non-cooperating and selfish users, each trying to optimize its own subjective goal. The focal points of this research are showing convergence to equilibrium point(s), showing the uniqueness of the equilibrium point, and analyzing the ability to influence the system into reaching a desired equilibrium.

The research community is currently working on lowering the key-management cost of secure multicast (cf. [Har99], [CGI+99]). In particular, Abdalla et al. [ASW99] suggested a key management scheme which occasionally tolerates a controlled number of free-riders. The authors showed that allowing some free-riders can substantially decrease the total number of keys each customer needs to store. Hence, from a technological point of view, allowing free-riders is advantageous when key storage is very limited, such as in applications using smartcards. In this paper we investigate, in a more general setting, the economic consequences that such an approach brings.

C. Our Approach

Let us imagine that a supplier of information services sells movies[1] to customers in the following manner. The supplier announces that at some particular time (9:00 pm Thursday) a particular movie (*Men in Black*) will be shown, and if you want to watch, you agree to pay a fixed subscription fee ($5). At 9 pm Thursday you turn on your TV set to the assigned channel and watch the movie.

As we discussed before, it is cheaper for the supplier to use broadcast than to use multicast. However, her revenue will dwindle to zero if she always uses broadcast. Thus a canny supplier will broadcast only some movies; naturally she will prefer to broadcast the ones for which large numbers of customers have paid, since those are the most expensive to multicast (and also the ones whose broadcast would provide free service to the fewest customers). In a full-information scenario (where, e.g., a customer can see how many of his colleagues are paying) this leads to an unstable multi-party game in which the players will be forced to adopt random strategies.

It is our contention that this will not happen in practice. There is almost certainly no way for a customer to find out what the other customers are doing, and there is more than adequate uncertainty in the evaluation of the popularity of a movie to prevent a customer from making a reliable judgment on the number of payers. Customers are not in a position to cooperate, and cannot be expected to do anything more than observe the supplier's behavior and act to maximize their individual expectations.

In what follows we propose a mathematical model for this system which is consistent with reasonable behavior, and permits analysis with as few arbitrary assumptions as possible. In the model a customer weighs his value v for seeing the movie and his estimate of the probability p of its being broadcast for free in order to decide whether he should pay. The value v is assigned randomly and depends on the "percentile popularity" of the movie; to compute p the customer combines his estimate (based on his own valuation v) of the movie's popularity and the supplier's past behavior. The supplier broadcasts all movies for which enough customers have paid, which will turn out to be the same as broadcasting all movies which are sufficiently popular.

The most arbitrary feature of the model is the choice of a parameterized family of probability densities describing how v varies among the customers for a movie of given popularity. We provide some guidance for selecting such families, and give some simple, reasonable examples which are then solved to determine the supplier's optimal policy.

[1] For ease of exposition, we will refer to the services being supplied as movies to be viewed on a TV. However, our model applies equally well to any other short-term subscription-based service.

II. The Model

A. The Customer

We begin by describing the customers' behavior. Fix a customer and a movie (that is, a particular showing of a movie); let v stand for the value attached by that customer to seeing the movie, expressed (say) in dollars. Thus if it were the case that he could see the movie only by paying the subscription fee λ, he would pay if $v > \lambda$ but not if $v < \lambda$. Since v will later be drawn from a continuous probability density, we do not need to concern ourselves with the zero-probability event that $v = \lambda$.

If the customer believes that with probability $p > 0$ the movie will be broadcast, so that he gets to see it for free, then he will of course be less willing to pay for it. Since his expected gain when he pays is $v - \lambda$ and when he doesn't, pv, he will pay when $v - \lambda > pv$ and not when $v - \lambda < pv$. This leads naturally to the definition of the *payment desirability* function

$$q(v) = 1 - \frac{\lambda}{v},$$

whereby the customer pays the subscription fee λ whenever his estimated broadcast probability p is less than $q(v)$.

We have assumed that the customer is rational and has a linear utility for money in the range of interest; neither will be precisely true but both are reasonable approximations in the aggregate.

The questions of how v is determined and how the customer computes p will be addressed shortly.

B. The Supplier

The behavior of the supplier is very simple indeed: she broadcasts all movies for which the fraction of paying customers exceeds some threshold B (which we call the *supplier's broadcast threshold*). We assume she will attempt to choose B so as to maximize her profit, and our model will in principle allow her to do that.

The result of her broadcast policy will be that a certain fraction b of movies will be multicast only to subscribers, and the rest will be broadcast to everyone; since the value b will be deducible by customers after sufficient time, their actions will be influenced by it. In particular, we should expect that a customer's estimate p of the probability that a movie will be broadcast will depend on b.

If individual customers *adapt* their strategies to the observed value of b, the aggregate effect will change b. Thus we are actually modelling a dynamical system, which has a feedback loop between the customers' strategies and b. We assume that after sufficient time, the system will stabilize to a fixed (pareto-optimal) point: The customers will be using strategies which will cause a fraction b of movies to be multicast, and for this particular value of b the customers would not have an incentive to change their strategies. We call the fixed point value b the *observed broadcast threshold*.

C. The Probability of Broadcast

A major issue is whether, and how, a customer's calculation of p should depend on the quality of the movie. The simplest assumption is of course that it doesn't; having observed that a fraction $1-b$ of the movies are broadcast, the customer merely takes $p = 1 - b$.

We believe, however, that customers will not fail to notice that it is the most popular movies which are broadcast. They will of course also be aware that opinions of movies (and time slots) vary, thus they can only guess at the popularity of a particular offer. Naturally that guess will, and should be, influenced by the customer's own valuation.

Thus we arrive at what we believe is the simplest and most natural way to model a non-trivial dependence of p upon b: the customer computes the *a posteriori* probability $p = p(b, v)$ that the movie will be broadcast, given that his own valuation of the movie is v. Of course, no customer will *literally* compute this value and most will not even think about the values v and b. But it is not unreasonable to expect a customer to think: "Hmm, I'd really like to see this one, but if I want to see it then it's likely others will also sign up, and I've noticed that a lot of the time popular movies are shown for free"... and the aggregate result ought to be something like the mathematical behavior of our model.

One weakness of this simple dependence is that it does not take into account the possibility that a customer wants to see a movie that he knows, perhaps from conversations with friends, is not popular. Another is that customers will have different ideas about the *a priori* distribution of movie valuations; however, even the assumption that p does not depend on v would require us to select a family of distributions for evaluations, so it costs us little to introduce this level of sophistication into the model.

D. The Rank of a Movie

To model the valuations, we introduce the concept of the *rank* r of a movie, as its percentile level of popularity (relative to all the movies offered by our suppplier). Thus a rank of $r = 3/4$ indicates a movie which is more popular than 75% of the movies shown; and the rank of a random movie is a uniformly random real from the unit interval $[0, 1]$.

Note that the notion of rank is in some sense external to the model. A priori, neither the supplier nor the customers actually know the rank of any particular movie with certainty. After the customers subscribe to a movie, the supplier may obtain a good estimate of the movie's rank, based on the number of subscribers—if a certain monotonicity assumption holds. We shall elaborate on this point shortly. The customers learn much less: in our model, the only *a posteriori* information they get about the movie's rank comes from the supplier's decision to use multicast or broadcast.

E. The Distribution of Valuations

We will assume throughout that both the number of customers and the (potential) number of movie offerings is large, and therefore we shall use continuous mathematics, and in particular probability density functions, to describe the distribution of valuations.

We assume that the valuations of a movie of rank r among the customers fit a probability density function f_r, so that the fraction of customers whose values fall between x and y is

$$\int_x^y f_r(v)dv \ .$$

This is again a simplification, since in reality two equally popular movies might have different valuation "profiles"—perhaps one has relatively uniform valuations while the other is highly controversial. Thus we must regard f_r as a kind of average density, and then it is reasonable to make the following "majorization" assumption:

> If $r' \geq r$ then the valuation density $f_{r'}$ *majorizes* f_r, that is, $\int_a^\infty f_{r'}(v)dv \ \geq \ \int_a^\infty f_r(v)dv$ for every a.

From a customer's point of view, the rank of a movie is *a priori* unknown, thus its valuation distribution is given by the mean density function

$$f(v) = \int_0^1 f_r(v)dr \ .$$

Later we will choose some particularly nice families of distributions on which to make test calculations. Until then we may represent various probabilities as integrals; for example, the probability that a random customer values a random movie higher than λ is $\int_0^1 \int_\lambda^\infty f_r(v)dvdr$. If no movies are broadcast, this quantity, multiplied by λ, would represent the mean revenue per customer per movie.

F. Monotonicity

We have said that our supplier will broadcast all movies paid for by at least a fraction B of the customers, which causes a fraction b of all movies to be *multi*cast, and the rest to be broadcast. To make the analysis tractable, we would like the following invariant to hold:

> The movies that are broadcast are precisely the movies of rank $r > b$.

To get this, we need the following, seemingly benign, yet critical, *monotonicity* property:

> The higher a movie is ranked, the more people will subscribe to it.

In view of our majorization assumption, it would suffice to show that an individual customer becomes more likely to pay for a movie if his valuation is higher; this certainly sounds reasonable. But it is also the case that in our model, as v goes up, so does the estimated broadcast probability p. Conceivably, the model might predict that a customer will cease wanting to pay for a movie when his valuation rises past some point, because he is then nearly certain that the movie is popular enough to be broadcast. We don't think this aberrant behavior is likely in practice. A customer's "natural monotonicity"—the positive correlation of his inclination to pay for an item, and his desire to have it—is not something which will be overcome by the soft information he has about the likelihood of broadcast. Hence, we will restrict ourselves to parameter ranges in which the monotonicity property holds.

Supposing monotonicity holds, a customer who has observed b and valued a particular offering at v will compute $p = p(b, v)$ as follows:

$$\begin{aligned} p &= \Pr(\text{movie broadcast} \mid \text{my valuation} = v) \\ &= \Pr(r > b \mid v) \\ &= \frac{\int_b^1 f_r(v)dr}{\int_0^1 f_r(v)dr} \end{aligned} \tag{1}$$

As we discussed in Section II-A, the customer will pay if $p < q$ (the payment desirability) and not if $p > q$. If monotonicity holds then we would expect there to be a single threshold $t = t(b)$ such that if $v > t$ then $p < q$ and if $v < t$ then $p > q$. We call this threshold t, if it exists, the *customer payment threshold*.

If a movie of rank r is offered, the fraction $F_{\lambda,b}(r)$ of the population that will pay for it is thus

$$\begin{aligned} F_{\lambda,b}(r) &= \Pr(\text{random customer has } v > t(b)) \\ &= \int_{t(b)}^\infty f_r(v)dv \ . \end{aligned} \tag{2}$$

F is subindexed by λ to remind the reader of the dependence on the subscription fee λ (through the customer payment threshold t).

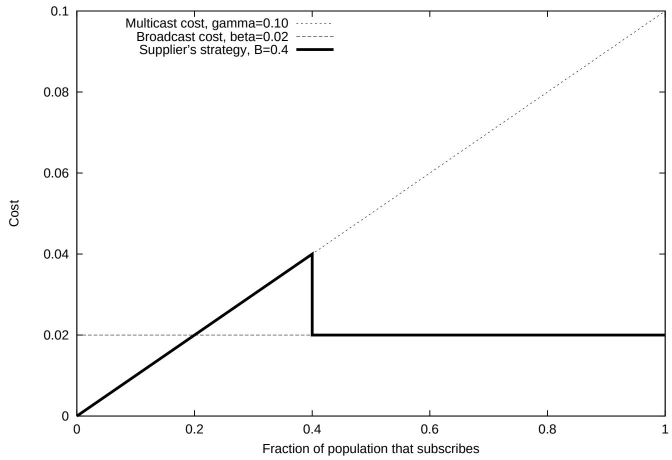

Fig. 1. The supplier's marginal cost structure.

v	The customer's valuation of a movie.
r	The rank of a movie.
$f_r(v)$	The probability density function of the valuations a movie of rank r.
B	The supplier's broadcast threshold.
b	The observed broadcast threshold.
$p(b,v)$	The customer's estimated broadcast probability function.
$q(v)$	The customer's payment desirability function.
t	The customer payment threshold.
λ	A movie's subscription fee.
γ	The per-customer multicast cost.
β	The per-customer broadcast cost.
$F_{\lambda,b}(r)$	The fraction of the population that pays for a movie of rank r.

TABLE I
The notation used in this paper.

The monotonicity property translates to an assurance that $F_{\lambda,b}(r)$ is monotone increasing in r. By the definition of B, the movies that are broadcast are those for which $F_{\lambda,b}(r) > B$, and since $F_{\lambda,b}(r)$ is monotone increasing this is equivalent to the condition $r > F_{\lambda,b}^{-1}(B)$. But by the definition of b, these movies coincide with the movies whose rank is $r > b$. Therefore, we conclude that the supplier's broadcast threshold is

$$B = F_{\lambda,b}(b) \ . \tag{3}$$

This formula allows us to optimize the observed broadcast threshold b in our analysis, and then to translate the best choice into a policy for the supplier.

G. Revenue, Cost, and Profit

According to the model, $F_{\lambda,b}(r)$ of the customers will pay for a movie of rank r, so we have, per customer:

$$\mathit{revenue}(r) = \lambda F_{\lambda,b}(r) \ . \tag{4}$$

Continuing to work on a per customer, per movie basis, we let γ be the marginal cost per customer for multicasting a movie only to its subscribers, and β the cost per customer for broadcasting to all of them. Thus if a fraction s of the customers pay for a movie, then the supplier's cost per customer is $s\gamma$ if the movie is multicast only to subscribers, and simply β if it is broadcast (see Fig. 1).

We assume that $\beta < \gamma < \lambda$, so that broadcasting is cheaper than multicasting to everyone, and the supplier is motivated, even while multicasting, to attract a maximum number of customers. This does not mean the supplier always makes a profit, because the fixed cost of offering a movie (which may well dominate the marginal costs) may put the supplier in the red. Here, however, we can without loss of generality ignore the fixed costs and define profit as if the fixed costs were zero.

According to the model, movies ranked above b are broadcast to everyone, and lower ranked movies are multicast only to their subscribers. Thus, the cost per customer per movie is

$$\mathit{cost}(r) = \begin{cases} \gamma F_{\lambda,b}(r), & r < b \text{ (multicast)}, \\ \beta, & r > b \text{ (broadcast)}. \end{cases} \tag{5}$$

Therefore, averaged over all movies,

$$\begin{aligned} \mathit{profit} &= \int_0^1 \mathit{revenue}(r)dr - \int_0^1 \mathit{cost}(r)dr \\ &= \lambda \int_0^1 F_{\lambda,b}(r)dr - \gamma \int_0^b F_{\lambda,b}(r)dr - \beta(1-b) \ . \end{aligned}$$

It remains only to choose b so as to maximize this quantity.

III. Examples

In this section we analyze in detail some examples of our model. Much of the notation we use is summarized in Table I for quick reference.

A. The Bilinear Density

Let us suppose that there is a natural maximum valuation, which we normalize to 1. We may imagine that for this unit price a customer buy the movie on a video cassette or DVD, thus no customer will value any movie higher than 1. Of course, after normalization we need $\lambda < 1$ to have any potential buyers.

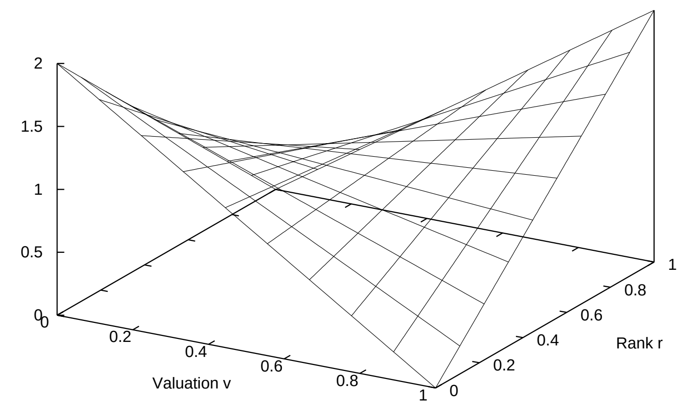

Fig. 2. The bilinear density $f_r(v) = 2rv + 2(1-r)(1-v)$.

In that case a natural *overall* valuation density (averaged over all ranks $0 \le r \le 1$) is the uniform density on the unit interval. The simplest non-trivial family of densities that enjoys a uniform overall valuation density is the family of the bilinear functions

$$f_r(v) = 2rv + 2(1-r)(1-v) \tag{6}$$

(see Fig. 2). A movie of rank $1/2$ has, like a random movie, a uniform valuation in $[0, 1]$.

B. An Individual Customer's Strategy

Assume that the supplier's broadcast threshold is set at B, and that the customers strategies have already stabilized to adapt to it. Thus, the observed broadcast threshold b is stable and known to the customers. The first step in our analysis is to compute a customer's *a posteriori* estimated broadcast probability $p = p(b, v)$, for a movie he values at v. Note that the bilinear density (6) is symmetric in v and r, so we have that $\int_0^1 f_r(v)dr = 1$ for any $v \in [0, 1]$. Plugging this into the definition of p in (1) we obtain that

$$p(b, v) = \int_b^1 f_r(v)dr = v(1-b^2) + (1-v)(1-b)^2 \ .$$

Fig. 3 depicts the resultant estimated broadcast probability p along with the payment desirability function $q = 1 - \lambda/v$. As we discussed in Section II-F, we need to ensure that the functions p and q have a single intersection point in the range $0 \le v \le 1$, at the customer payment threshold $v = t(b)$. For this to hold, it suffices to ensure that $p(b, 1) < q(1)$, since $p(b, v)$ is linear in v, $q(v)$ is concave, and $p(b, \lambda) > q(\lambda) = 0$ for any b. Therefore, monotonicity holds if and only if

$$b > \sqrt{\lambda} \ . \tag{7}$$

Once we are assured that a single intersection point exists, we can compute the customer payment threshold $t(b)$

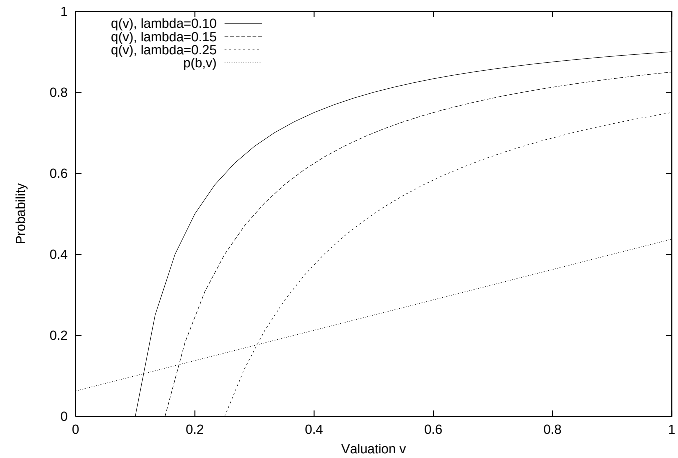

Fig. 3. The estimated broadcast probability $p = p(b, v)$, and the payment desirability $q = q(v)$, as functions of the valuation v, for an observed broadcast threshold of $b = 0.75$ and different values of λ.

by solving the equation $p(b, v) = q(v)$. In our case $t(b)$ is the smaller of the two roots of the quadratic equation

$$v^2(2b - 2b^2) + v(b^2 - 2b) + \lambda = 0 \ .$$

The customer payment threshold t completely characterizes the strategy of individual customers: A customer would pay for any movie which he values at $v > t$.

Remark: Requiring that b is bound from below is a reasonable condition. For instance, it is not hard to check that if $b < \lambda < 3/4$ then $p(b, v)$ and $q(v)$ do not intersect at all. In other words, when the observed broadcast threshold b is so low, no customer would ever pay for any movie. But this is behavior cannot be sustained more than briefly, since when no customer pays, no movies are broadcast, leading to a new observed broadcast threshold of $b = 1$. Thus very small values of b cannot be stable points in our game.

C. The Population's Aggregate Behavior

After we know the customer payment threshold t, we can turn to analyzing the aggregate behavior of the customer population. We calculate the fraction of the population that pays for a movie of rank r by plugging (6) into (2) to obtain

$$F_{\lambda,b}(r) = \int_t^1 f_r(v)dv = r(1-t^2) + (1-r)(1-t)^2 \ .$$

By (3) the supplier's broadcast threshold is $B = F_{\lambda,b}(b)$, and since $F_{\lambda,b}(r)$ is linear in r, we can invert it to obtain a simple formula for the observed broadcast threshold b as a

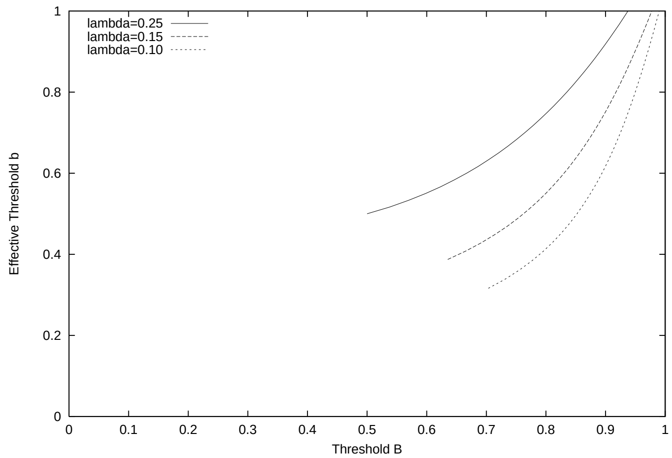

Fig. 4. The observed broadcast threshold b as a function of the supplier's broadcast threshold B, for different values of λ.

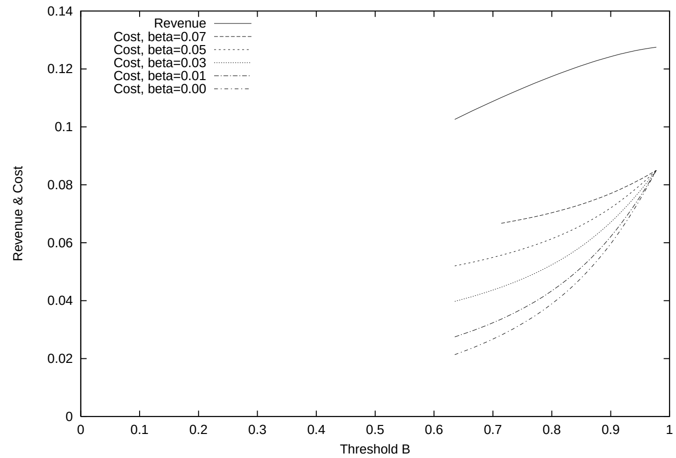

Fig. 5. The revenue and cost as a function of the supplier's broadcast threshold B, for $\lambda = 0.15$, per-customer multicast cost $\gamma = 0.10$, and different per-customer broadcast costs β.

function of B:

$$b = \frac{B - (1-t)^2}{2t(1-t)} \ .$$

Fig. 4 depicts b as a function of B for the bilinear density (b also depends on λ, via the customer payment threshold t). The curves are not defined for all values of B because of requirement (7); for instance, the curve for $\lambda = 0.15$ is defined only for $B \geq 0.635$, which corresponds to $b \geq 0.387 = \sqrt{0.15}$. Lower values of B, which imply lower values of b, would lead to non-monotonic customer behavior. We can also see that b reaches 1 (i.e., no movies are ever broadcast) for B values lower than 1, e.g., when $\lambda = 0.15$ we get $b = 1$ for $B = 0.9775$. This simply means that at most 97.75% of the population ever subscribes to a movie.

D. Revenue, Cost, and Profit

From (4) we have that the aggregate revenue is

$$revenue = \lambda \int_0^1 F_{\lambda,b}(r)dr = \lambda(1-t) \ .$$

This can also be seen directly from the fact that the overall density is uniform, so if users pay only for movies valued at $v > t$, the fraction of the population that subscribes to the average movie is $1 - t$.

Recall that γ is the per-customer multicast cost, and that the per-customer broadcast cost is $\beta < \gamma$. From (5) we have that the aggregate cost is

$$\begin{aligned} cost &= \gamma \int_0^b F_{\lambda,b}(r)dr + \beta(1-b) \\ &= \gamma b[1-(2-b)t+(1-b)t^2] + \beta(1-b). \end{aligned}$$

Fig. 5 depicts the revenue and cost curves, as functions of the supplier's broadcast threshold B. The curves are only defined in the range of B values for which $\sqrt{\lambda} < b < 1$ (e.g., $0.635 < B < 0.9775$ for $\lambda = 0.15$). However, there is an additional restriction on B: If the fraction s of subscribers is small, the cost of multicasting to them, γs, may be lower than the broadcast cost β. So we can limit ourselves to $B > \beta/\gamma$. If β is close enough to γ, this bound is more restrictive than the bound implied by $b > \sqrt{\lambda}$. This explains why the curve for $\beta = 0.07$ starts at $B = 0.715$, rather than at $B = 0.635$.

We believe that typically $\beta \ll \gamma$, since in the model these values are *per customer*. In reality, the cost of broadcasting is often independent of the subscriber population size (the cost of a satellite transmission is independent of the number of receivers), so our β would actually decrease as the total number of subscribers grows. In particular, the value $\beta = 0$ may well be reasonable, for instance, if a unicast to a single subscriber has the same cost as a broadcast to everyone—which is the case for satellite communication.

Fig. 6 depicts the supplier's profit. We see that for all the parameter settings we tried, the profit is maximized if B is set to a value between 0.86 and 0.69. The cheaper it is to broadcast, the more pronounced is the increase in profit. The most striking case is when broadcast is "for free" ($\beta = 0$): If the supplier broadcasts whenever more than 0.69% of the population subscribes to a movie, her per-customer profit will be 93% higher than if she never broadcasts ($B = 1$).

We have experimented with other settings of λ, γ, and β for the bilinear density function. The results are qual-

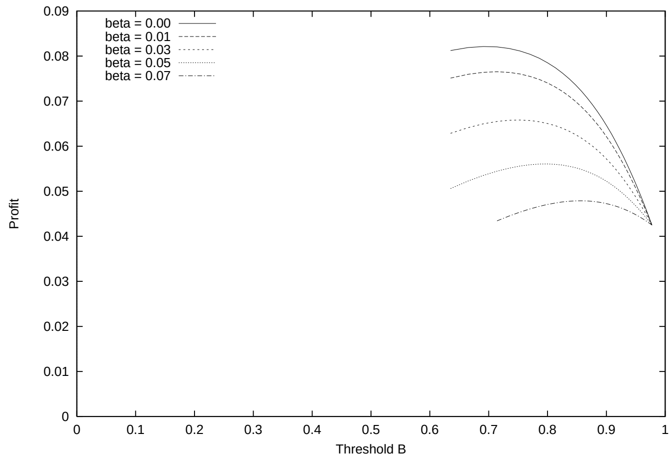

Fig. 6. The profit as a function of the supplier's broadcast threshold B, for $\lambda = 0.15$, per-customer multicast cost $\gamma = 0.10$, and different per-customer broadcast costs β.

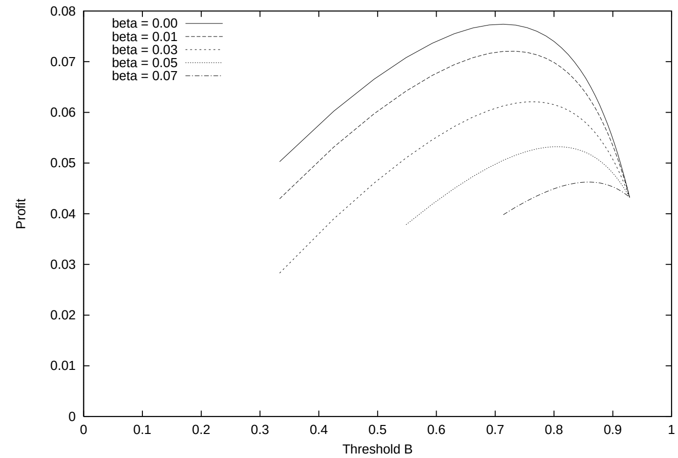

Fig. 7. The profit as a function of the supplier's broadcast threshold B, for the parabolic density, with $\lambda = 0.15$, per-customer multicast cost $\gamma = 0.10$.

itatively the same, with different numerical values, so we omit the graphs.

E. Other Density Functions

In order to test the robustness of our model, and to estimate the extent to which the conclusions are artifacts of the choice of density function, we experimented with some other families of density functions.

E.1 Parabolic Density

In this family of densities, the density function of a movie of rank r is

$$f_r(v) = -v^2 + 2rv + 4/3 - r \ .$$

These densities have the property that the mode (= peak) of $f_r(v)$ is at $v = r$, for all $r \in [0, 1]$. Their overall density, for an average r, is $f(v) = -v^2 + v + 5/6$, i.e., a non-uniform overall density.

Following the same sequence of computations we used for the bilinear density we obtain that

$$p(b, v) = \frac{v^2(b-1) + v(1-b^2) + 5/6 - 4b/3 + b^2/2}{-v^2 + v + 5/6} \ .$$

The customer payment threshold t is a solution to the equation $p(b, v) = q(v)$, which is the cubic equation

$$g(v) = v^3 b - v^2(b^2 + \lambda) + v(b^2/2 - 4b/3 + \lambda) + 5\lambda/6 = 0 \ .$$

The leading coefficient of the cubic function g is positive, and $g(0) = 5\lambda/6 > 0$. Therefore, ensuring that $g(1) < 0$ suffices to prove that a single solution to $g(v) = 0$ exists in the range $[0, 1]$, and therefore to guarantee the monotonicity property. The condition $g(1) < 0$ translates to ensuring that b is not too small:

$$b > \frac{-1 + \sqrt{1 + 15\lambda}}{3} \ .$$

Once we know the customer payment threshold t, it is easy to see that

$$F_{\lambda,b}(r) = \int_t^1 f_r(v)dv = t^3/3 - rt^2 + (r - 4/3)t + 1.$$

The resultant profit curves appear in Fig. 7, using the same parameter settings we used in Fig. 6. For the parabolic density, we see that if broadcast is "for free" ($\beta = 0$), and the supplier's broadcast threshold is set at $B = 0.714$, then her profit would be 79% higher than had she never used broadcast.

E.2 Cubic Density

In this family of densities, the density function of a movie of rank r is

$$f_r(v) = (4 - 8r)v^3 - (6 - 12r)v^2 + (2 - 2r).$$

These densities have a derivative $\frac{\partial f_r(v)}{\partial v} = 0$ at $v = 0$ and $v = 1$ for all $r \in [0, 1]$, and the overall density, for an average r, is uniform. Furthermore, $\int_0^1 f_r(v)dr = 1$ for all $v \in [0, 1]$. Straightforward calculations show that

$$p(b, v) = (4b^2 - 4b)v^3 + (6b - 6b^2)v^2 + 1 - 2b + b^2 \ ,$$

and that the customer payment threshold t is a solution to the quartic equation

$$(4b - 4b^2)v^4 + (6b^2 - 6b)v^3 + (2b - b^2)v - \lambda = 0 \ .$$

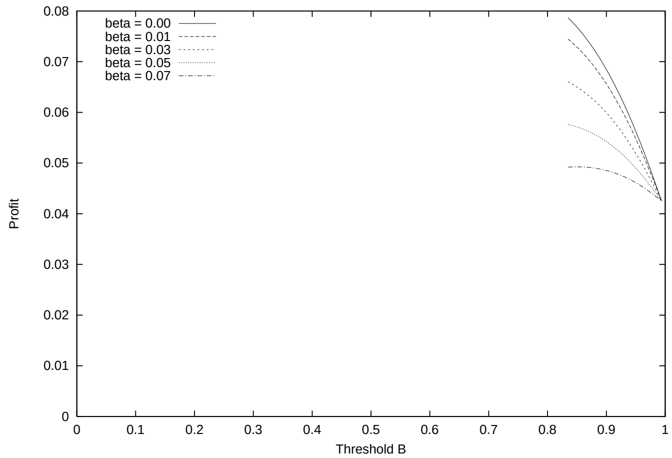

Fig. 8. The profit as a function of the supplier's broadcast threshold B, for the cubic density, with $\lambda = 0.15$, per-customer multicast cost $\gamma = 0.10$.

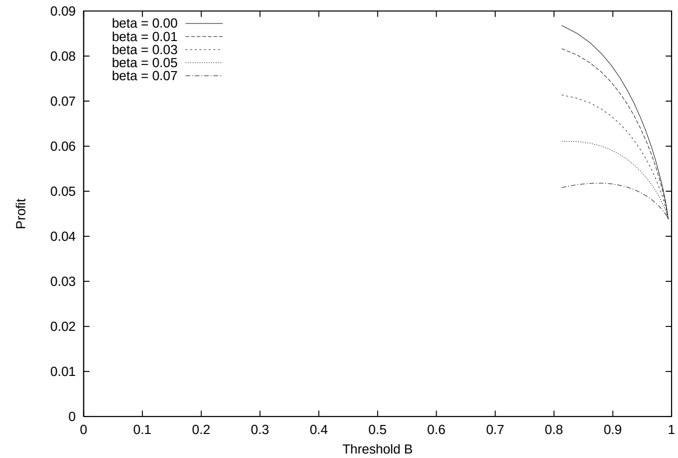

Fig. 9. The profit as a function of the supplier's broadcast threshold B, for the cosine density, with $\lambda = 0.15$, per-customer multicast cost $\gamma = 0.10$.

Requiring that $b > \sqrt{\lambda}$ ensures that $q(1) > p(b, 1)$, which is a necessary condition for monotonicity. However, for the cubic density this condition is not sufficient, since in principal p and q may intersect in up to 4 points in $[0, 1]$. But it is not hard to show that if $b > 11/19$ then only one intersection point occurs in $[0, 1]$. Therefore, if $b > \max\{\sqrt{\lambda}, 11/19\}$ the monotonicity is guaranteed.

The profit curves are shown in Figure 8. Qualitatively, we see the same picture we have seen for other density functions: Our model predicts that setting the supplier's broadcast threshold at $B = 0.84$ would lead to an 85% increase in the supplier's profit, in comparison to the "never broadcast" strategy. The graphs seem to indicate that the profit would be even larger for values of $B < 0.84$. However, such low supplier's broadcast thresholds would lead to non-monotonic behavior, and hence cannot be modeled in our current framework.

E.3 Cosine Density

We have also experimented with a family of density functions that are based on trigonometric functions, which we call the *cosine density*. The density function of a movie of rank r is "half the period" of a cosine curve, phase-shifted so the maximum is at $v = r$, and normalized to give an integral $\int_0^1 f_r(v)dv = 1$. Formally,

$$f_r(v) = \frac{\pi(1 + \cos(\pi(v - r)))}{\pi + 2\sin(\pi r)} .$$

For this family of densities the expressions for the various components of our model are quite unwieldy, so we omit all the details. The final profit curves are shown in Fig. 9, and again, they show the same behavior we have seen all along, with slightly different numerical values.

IV. Conclusions

In this work we built an economic model for multicast services. A central component in our model is that when enough customers subscribe to a service, the supplier broadcasts the service to all the customers. From this model, coupled with some mild assumptions on the supplier's cost structure, we can find the optimal setting for the supplier's supplier's broadcast threshold, and the optimal strategy for rational customers.

In all the examples we studied, our model predicts that the supplier's profits will be maximized if the supplier's broadcast threshold is set below 100%: The loss in revenue due to customers subscribing to fewer services is offset by the cost savings made possible by broadcasting the most popular services to all customers.

In conclusion, we find that our model is analyzable, reasonably consistent with expected customer behavior, and fairly robust with respect to parameter choices. As such, we believe it can be of value to a supplier in devising a multicast/broadcast strategy, and that broadcasting when subscriptions are sufficiently high is likely to be the approach of choice in maximizing profits.

References

[AB98] E. Altman and T. Basar. Multiuser rate-based flow control. *IEEE Transactions on Communications*, 46(7):940 – 949, July 1998.

[AFM92] S. Armstrong, A. Freier, and K. Marzullo. Multicast transport protocol, February 1992. Internet RFC 1301.

[ASW99] M. Abdalla, Y. Shavitt, and A. Wool. Towards making broadcast encryption practical. In *Financial Cryp-*

tography'99, Anguilla, BWI, February 1999. LNCS 1648, Springer-Verlag.

[BMS99] D. Balenson, D. McGrew, and A. Sherman. Key management for large dynamic groups: One-way function trees and amortized initialization. Internet Draft, February 1999.

[CGI+99] Ran Canetti, Juan Garay, Gene Itkis, Daniele Micciancio, Moni Naor, and Benny Pinkas. Multicast security: A taxonomy and efficient constructions. In *IEEE INFOCOM'99*, March 1999.

[DC85] S. E. Deering and D. T. Cheriton. Host groups: A multicast extension to the internet protocol, December 1985. Internet RFC 966, Stanford University.

[Har99] T. Hardjono. Secure multicast group (SMuG) reference framework. `http://www.ipmulticast.com/community/smug`, March 1999. IRTF Work in Progress.

[HSE97] S. Herzog, S. Shenker, and D. Estrin. Sharing multicast costs. In L. W. McKnight and J. P. Bailey, editors, *Internet Economics*, pages 169–212. MIT Press, Cambridge, Mass., 1997.

[KL95] Yannis A. Korilis and Aurel A. Lazar. On the existence of equilibria in noncooperative optimal flow control. *J. ACM*, 42(3):584–613, May 1995.

[KO99] Yannis A. Korilis and Ariel Orda. Incentive-compatible pricing strategies for QoS routing. In *IEEE INFOCOM'99*, March 1999.

[KVA98] Yannis A. Korilis, T. A. Varvarigou, and Sid R. Ahuja. Incentive-compatible pricing strategies in noncooperative networks. In *IEEE INFOCOM'98*, pages 439 – 446, March 1998.

[LO97] L. Libman and Ariel Orda. Atomic resource sharing in noncooperative networks. In *IEEE INFOCOM'97*, pages 1006 – 1013, April 1997.

[MB97] L. W. McKnight and J. P. Bailey, editors. *Internet Economics*. MIT Press, Cambridge, Mass., 1997.

[MQ95] B. M. Macq and J.-J. Quisquater. Cryptology for digital TV broadcasting. *Proceedings of the IEEE*, 83(6):944–957, 1995.

[She95] S. J. Shenker. Making greed work in networks: a game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking*, 3(6):819 – 831, December 1995.

[WGL98] C. K. Wong, M. Gouda, and S. S. Lam. Secure group communications using key graphs. In *Proc. ACM SIGCOMM*, pages 68–79, Vancouver, BC, 1998.

[WHA98] D. M. Wallner, E. J. Harder, and R. C. Agee. Key management for multicast: Issues and architectures. Internet Draft, September 1998. Available from `http://www.ietf.org/ID.html`.