A Profitable Multicast Business Model

Sneha Kumar Kasera, Markus Hofmann
Bell Labs Research

Raymond E. Miller
University of Maryland at College Park

I. Introduction

Multicast is an efficient paradigm for transmitting data from a sender to a group of receivers. Multicast incurs lower network bandwidth and end–system costs than broadcast to all receivers or multiple unicasts to individual receivers. One of the main impediments in wide scale deployment of multicast is lack of a good business model. Any technology needs a good business model to succeed. Each of the parties involved must see some advantage in using the technology. In this paper we present a simple business model for multicast in the Internet that uses the inherent benefits of multicast to make it profitable to all the parties, including the multicast sender, multicast receivers and the network providers, that are involved.

Our model is based on the following principle. *Multicast receivers should not pay any extra charge for receiving multicast over unicast, the sender pays for the bandwidth used in multicast to the Internet service providers and might charge the receivers for the content*. We describe our model and demonstrate that its use proves profitable to the sender, receivers and the Internet service providers. We also discuss issues related to this model's implementation.

II. Related Work

Multicast pricing is receiving increasing attention in the research community. Several interesting ideas [1], [3] have been proposed. Unfortunately none of the ideas have been able to fully motivate all the parties involved in multicast including the sender, the network providers and the receivers. Chuang and Sirbu [2], through extensive simulations of over a wide range of networks, have shown that the ratio of links in the multicast tree from a sender to n receiver sites (or POPs) to the average number of links in unicast paths from the sender to n receiver sites is $n^{0.8}$. Our work differs from Chuang and Sirbu's work in the following significant ways. First, we propose how the sender charge could be divided among the network providers. Second our business model clearly identifies the benefits of the sender, the receivers and the network providers. In our numerical examples, we also exploit the region between $n^{0.8}$ and n to provide the benefits of multicast to all parties.

III. Business Model

Multicast is inherently beneficial. A good multicast business model should be able to provide incentives for all the parties involved. In this section we present our business model that meets this requirement. In presenting our business model we use a scenario where a sender is multicasting multimedia data to a group of N receivers, however we believe that our approach generalizes to other scenarios, as well. The design principles of our business model are as follows:

- receivers do not pay any extra for bandwidth for receiving multicast versus unicast
- the sender pays for the multicast bandwidth
- the sender might charge the receivers for the content
- the network providers should receive revenue based on the proportion of their resources that are used in multicast.

Let us now see how each of the parties involved in the multicast, the sender, receivers and the ISP(s), can benefit by using multicast. As far as a receiver is concerned it does not matter whether it is receiving data through unicast or multicast. Typically a receiver pays a fixed fee to its ISP. There is no incentive for the receiver to pay or share the cost of multicast. A receiver might be interested in receiving multicast if the content is offered free or at a discounted rate when multicast. A sender will be interested in multicast if it makes sure that its data reaches N receivers and by using multicast it pays less than it would for N unicast connections. It can then use a part of this profit to reduce the price of the content. It will be argued later that as the number of receivers increases the reduction in cost due to using multicast over unicast also increases. Hence the sender can afford to increase the discount of the content as the number of receivers increase. This strategy is similar to the one used in many online stores that reduce the price of an item depending on the amount sold. A sender could also multicast free content for an Internet radio–like service where it covers the cost of multicast as well as makes profit through commercial advertisements.

There are several reasons why a network provider (Internet service provider or ISP, National service provider or NSP) might be interested in providing multicast. First, a service provider uses the same or less bandwidth for multicasting to N receivers than individually unicasting to each receiver. Hence it could offer multicast service at a cost that is less than the cost of N times the unicast cost. One might argue that this will lead to a reduction in revenues for the ISP because it is reducing its business from N unicast connections to one multicast "connection" that is offered at a lower cost. The argument in favor of multicast is that an ISP can accommodate only a certain number of unicast connections across a bottleneck link whereas it can support an equal number of multicast connections through that link. By charging incrementally more for multicast it can serve more customers and get more revenue. In fact, our model shows that an ISP could charge more for a certain bandwidth in multicast than in unicast and still leave room for sender profit. An ISP might also need to provide multicast service to stay competitive. In the next section we analyze the profits of the sender, receivers and the ISP when they use multicast.

IV. Profit Analysis

The symbols used in our analysis are given in Table I. The sender profit due to choosing multicast over unicast can be expressed as follows:

$$X = (Nx_m - C) - (Nx_u - Nb_u) \qquad (1)$$

The first term is the sender profit when it uses multicast and the second term is the sender profit when it uses unicast. The

x_u	–	cost of content per receiver when unicast
x_m	–	cost of content per receiver when multicast
b_u	–	sender bandwidth charge of unicast to a single receiver
N	–	number of receivers
n	–	number of receiver sites or POPs
X	–	sender profit (or loss) due to using multicast over unicast
Y	–	receiver savings due to using multicast over unicast
Z	–	network provider profit
C	–	the multicast bandwidth usage charge

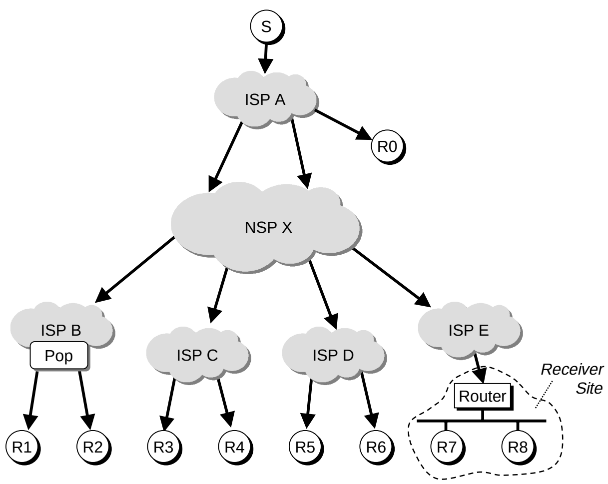

Fig. 1. Multicast Involving Multiple ISPs

expression for C is derived in this section.

The receiver savings can be simply expressed as follows:

$$Y = x_u - x_m \tag{2}$$

Multicast to receivers scattered over wide area networks might involve several ISPs (as well as NSPs). A typical scenario (shown in Figure 1) involves the sender ISP A, an NSP X providing backbone access and several ISPs A–E, serving the receivers (R0–R8), connected to this NSP. There are two important problems. First, how much do the ISPs charge the sender for multicast and second how is the charge paid by the sender shared among the multiple ISPs. In the remainder of this section we present solutions to both these problems.

We first determine the charge of an individual ISP. The part of the multicast tree inside an ISP can have multiple input and output links. For example, the backbone NSP X in Figure 1 has two input and four output links. Multiple input links to an ISP suggests that the ISP will have multiple subtrees of the multicast tree. We now derive an expression for the number of links in potentially multiple subtrees of the multicast tree inside an ISP as a function of the number of input and output links. Let I be the number of input links and let J be the number of output links. Let K be the branching factor of each subtree of the multicast tree in the ISP. We start with the assumption that K is an

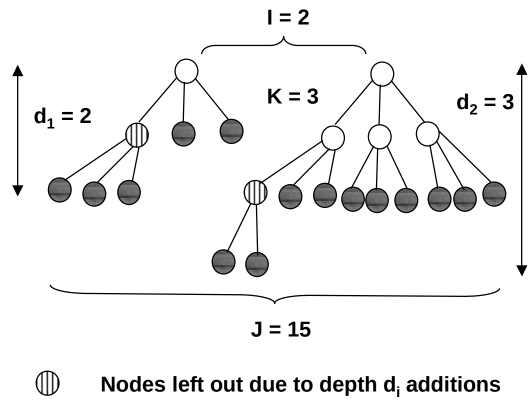

Fig. 2. Two subtrees inside an ISP

integer greater than one. We also assume that the subtrees are constructed in such a way that each of the internal nodes has K branches except those nodes that are parents of leaf nodes (see Figure 2). Depending on the value of J not all nodes that are parents of the leaf nodes of the subtrees will have K branches. We will find that the expression we derive holds, even when K represents the average branching factor, which is not necessarily an integer, for any general subtrees.

Let L_m be the number of multicast links in the subtrees of the multicast tree in an ISP p. Then

$$\begin{aligned} L_m \;=\; & \sum_{i=1}^{I} \sum_{j=1}^{d_i-1} (K^j + (J - \sum_{i=1}^{I} K^{d_i-1}) \\ & + \lceil (J - \sum_{i=1}^{I} K^{d_i-1})/(K-1) \rceil) \end{aligned} \tag{3}$$

The first term in Equation (3) is the number of links in I K-ary trees of depth $d_i - 1$ (where $i = 1, 2, 3, \ldots, I$). The second term is the sum of the additional links needed at depth d_i in all trees. The third term is the sum of the number of additional links needed in each tree i due to leaves of depth $d_i - 1$ being eliminated by depth d_i additions. By using simple algebra Equation (3) can be reduced to

$$L_m = \lceil (K(J - I))/(K-1) \rceil \tag{4}$$

We now prove that the above expression is true for general trees. We redefine K to be the average branching factor of I general trees. Now the average branching factor, and hence K, is the ratio of L_m and the number of internal nodes in the I trees. Noting that $(J - I)$ is L_m minus the number of internal nodes in the I subtrees, the right hand side of the above expression reduces to L_m. This shows that we can easily find the number of links in the portions of the multicast tree inside an ISP as a function of the number of input and output links and the average branching factor. Equation (4) is true only when $K > 1$ and when $J \neq I$. When $K = 1$ or when $J = I$, we set $L_m = I$.

If d is the average length of unicast paths in the subtrees of the multicast tree in the ISP p and if c_u is the cost of unicast across it[1], then an ISP p can charge C_p by using the formula

$$C_p = L_m * c_u / d \tag{5}$$

[1] Note that c_u is different from b_u. c_u is cost of unicast across an ISP and b_u is the charge that the sender pays for unicast to its ISP which in turn might be paying a portion of this charge to other neighboring ISPs.

An ISP could incrementally charge more than C_p to make profit (and also to account for additional costs in constructing multicast trees and keeping multicast state at routers). Hence an ISP could potentially charge $C_p(1+\delta)$ where $\delta > 0$. The profit of the network provider p is given by the following expression.

$$Z = C_p \delta \tag{6}$$

We propose that each receiver ISP can compute its share and pass it to the next ISP above it in the multicast tree towards the sender which adds its own share to the cost received from down-tree ISPs and sends it up the tree towards the sender. Finally the sender receives the total cost that it has to pay. If there are S ISPs then the sender has to pay $C = (1+\delta)\sum_{p=0}^{S} C_p$. It should be noted that this amount should be sufficiently less than Nb_u otherwise the sender will not have much incentive for using multicast.

V. Numerical Examples

We now present some numerical examples to demonstrate the profitability of our business model. For simplicity, we will only consider the case when the sender as well as the receivers are served by the same ISP (say ISP 1). In the single ISP case, for the purpose of constructing realistic examples, we use the Chuang and Sirbu law [2] to find C_1. In [2], Chuang and Sirbu have shown that for multicast trees constructed from a wide range of topologies the ratio of links in the multicast tree from a sender to n receiver sites (POPs) to the average number of links in unicast paths from the sender to n receiver sites is $n^{0.8}$. Using the Chuang and Sirbu law, L_m/d in Equation (5) can be replaced by $n^{0.8}$. Also note that the charge of unicast across the ISP c_u now becomes b_u for the single ISP case. The sender charge could be expressed as $n^{0.8+\epsilon} b_u$. We add ϵ to the exponent to account for the factor δ. δ should be chosen such that ϵ lies in $(0, 0.2)$ else the sender charge for multicast to n receiver sites will be more than the charge of n unicasts. The sender and ISP profit will depend upon the choice of δ (or ϵ). The relation between ϵ and δ can be expressed as $(1+\delta) = n^{\epsilon}$.

We consider a scenario where the sender charge by the network provider for unicast to one receiver, b_u, is 10. The amount the sender charges the receiver when it unicasts the content, x_u, is 15. We also assume that $N = n$, i.e., there is only one receiver per receiver site. As the charge paid by the sender, to the network provider for multicast, increases only with increase in n, the sender profit determined below is the minimum sender profit.

Figure 3 shows how the sender profit due to multicast increases with the increase in number of receiver sites for different values of δ. Figure 4 shows how the network provider profit increases with the increase in number of receiver sites for different values of δ. In both these figures, the discount offered by the sender to a multicast receiver, $(x_u - x_m)/x_u$, is 25%. Note that higher δ means higher network provider profit and lower sender profit. When the number of receivers is small ($n <= 10$) then the sender incurs a loss even for $\delta = 0.10$. As the number of receivers increases the benefits of multicast over unicast increases and this behavior is reflected in the increased sender and network provider profit.

Figure 5 shows how the sender profit due to multicast increases with the increase in the number of receiver sites for different values of the discount offered by the sender to receiver for using multicast. In this figure $\delta = 0.30$. We see that as the number of receiver sites increases the sender is able to afford higher discounts.

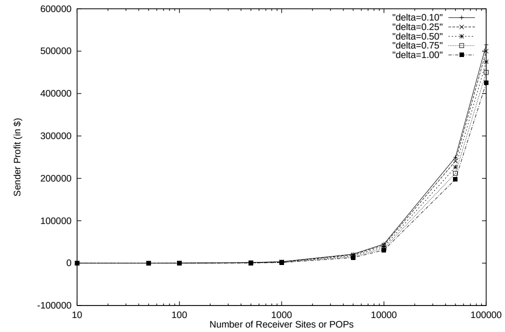

Fig. 3. Sender Profit when Receiver Discount=25%

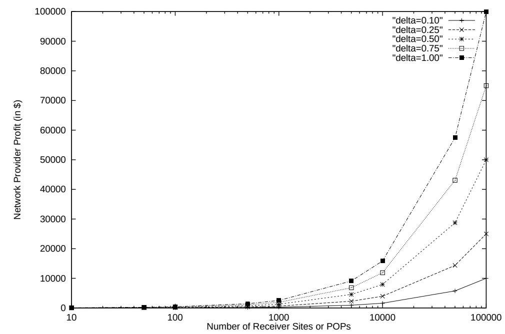

Fig. 4. Network Provider Profit

VI. Deployment Issues

Although the proposed business model and pricing scheme is independent from actual implementation techniques, there are a few implications for its deployment in real networks. This section will discuss how various parameters in our business model can be determined in practice, how ISPs will interact with each other, what their relationship will be, and opportunities for new services driven by our multicast business model.

Our scheme allows ISPs to calculate their fair share of the sender charge based on their local view and knowledge of the multicast transmission. It is not necessary for the ISPs to know the shape of the entire multicast tree, but just the section in their

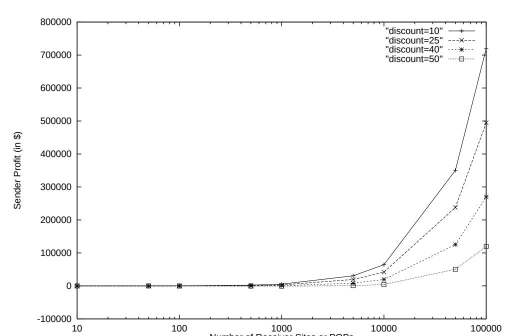

Fig. 5. Sender Profit when $\delta = 0.30$

own domain. In particular, it is sufficient for an ISP to consider the number of links coming into its domain (I) and the number of outgoing links (J) per multicast session. This information can easily be obtained at the border routers of the ISP. It is not necessary to do any monitoring or to maintain any state in the core of the network, nor is it necessary to introduce additional signaling or control protocols. This improves scalability and allows for easy deployment without additional complexity. Besides these two session–specific parameters, each ISP has to determine the average branching factor (K) and the average depth (d) of the multicast subtree in its domain. We envision that ISPs will start monitoring these parameters on a per-session basis during an initial deployment phase, but will gain enough experience to determine flat values based on the characteristics of their network. The average depth (d) of the multicast subtree, for example, could be set to the average length of a unicast path through the ISP's network. Similarly, the ISP might use the average fan-out of multicast routers in its network as the average branching factor (K). Network specific simulations and real-life experience will provide more insights on how to dimension these parameters. Practical experience will also provide feedback on meaningful values for the profit factor δ from a business and a technical perspective.

Using the local parameters mentioned above, each ISP calculates its fair share in the sender's charge. Shares from different ISPs will then be accumulated in a bottom up manner, i.e. they will be accumulated from receiver ISPs (leaf ISPs) up to the sender ISP (root ISP). In Figure 1, for example, each leaf ISP (B, C, D and E) calculates its individual multicast charge and gets paid for it by the parent NSP X. NSP X accumulates the amount paid to its child ISPs (B, C, D and E), adds its own multicast charge and gets paid for the accumulated amount by its parent ISP A. Finally, ISP A charges the sender for the entire amount, including the individual cost of ISP A. It is up to the sender to charge receivers for delivered content, thus covering part of the transmission cost. The sender also decides whether parts of its savings will be passed on to receivers in form of a discount.

The proposed pricing scheme requires interaction only between neighbor ISPs, there is no need for information exchange between non-neighbor ISPs. This reflects existing (bilateral) business relationships and avoids the need for additional partnerships. However, our model presupposes a trust relationship between the ISPs, because it implicitly assumes that each ISP is using reasonable values for K and d. This might not be taken for granted, because manipulating these values allows ISPs to gain more profit. While actual verification and security mechanisms are out of the scope of this paper, our approach does include hooks for integration of such features. For example, ISPs could include flat values for K and d into bilateral service level agreements with neighbor ISPs. This allows them to verify the multicast charge of their neighbors, assuming that ISPs will report the number of output links for a multicast session to their parents. Together with signed charge reports that are passed along the ISP tree towards the root ISP, this feature allows the sender to verify the correctness of a multicast charge. Other mechanisms may exist and details need to be worked out. The point is, however, that our business model provides a reasonable and simple mechanism for verification purposes.

As we have shown in Section V, our business model allows a sender to give higher discounts as the number of receiver sites increases. This opens an opportunity for a new breed of content delivery services. Content providers could offer delivery of specific content at certain points in time. Interested customers would sign up for a specific delivery time and would see the price decrease as more and more customers sign up. For example, a movie could be offered for delivery in one hour, two hours, three hours, etc.. The longer delivery time a customer accepts, the higher the chances that more customers sign up and that the price for the movie drops. This model is similar to the purchasing service offered by some online stores, where the price of a product drops as the number of buyers increases. Our business model enables a variety of similar services in the context of content delivery.

VII. Conclusions

One of the main reason for multicast technology not being deployed on a wide scale is the lack of a good business model. We have proposed a simple business model that uses the inherent benefits of multicasting to make it attractive and profitable to all parties, including the sender, the receivers, and the network providers. In our scheme, the responsibility for paying for multicast transmission lies with the sender. Each ISP will calculate its "fair" share of the sender charge based on its local view and knowledge. Receivers will not be charged any additional fee for receiving multicast data. While this principle makes sense from a business point of view, it opens technical challenges that need to be solved. Mechanisms need to be in place for disabling receivers from joining an arbitrary multicast session, thus increasing the sender's cost. Also, receivers should be prevented from joining non–existent multicast sessions to avoid unnecessary ISP cost. Other issues for future work include verification and security issues, as well as the trust model between ISPs.

References

[1] B. Briscoe, *The Direction of Value in Connectionless Networks*. Proceedings of 1st International workshop on Networked group Communication (NGC), November 1999.

[2] J. Chuang and M. Sirbu, *Pricing Multicast communications*. In Proceedings of the INET'98, 1998.

[3] S. Herzog, S. Shenker and D. Estrin, *Sharing the "Cost" of Multicast Trees: An Axiomatic Analysis*. Proceedings of ACM SIGCOMM, August 1995.