

## Fair Allocation of Multicast Transmission Costs

Patrik Österberg and Tingting Zhang  
Department of Information Technology and Media  
Mid Sweden University  
SE-851 70 Sundsvall, Sweden  
patrik.osterberg@miun.se, tingting.zhang@miun.se

### Abstract

*In scenarios where many receivers simultaneously are interested in the same data, multicast transmission is more bandwidth efficient than unicast. The reason is that the receivers of a multicast session share the resources through a common transmission tree. Since the resources are shared between the receivers, it is reasonable that the costs corresponding to these resources should be shared as well.*

*This paper deals with fair cost sharing among multicast receivers, and the work is based upon the assumption that costs should be shared according to the resource usage. However, it is not for certain that an optimally fair cost allocation is most beneficial for the receivers; receivers that cannot cover their fair share of the costs may nevertheless be able to contribute to the cost sharing to some extent. We propose a cost-allocation mechanism that strives to allocate the costs fairly, but gives discount to poor receivers who at least manage to cover the additional cost of providing them with the service.*

**Keywords:** *multicast, fairness, cost allocation*

### 1. Introduction

Video-streaming services are rapidly gaining in popularity, and the quality of these services is also increasing. Internet video already has attracted a large crowd, but the quality leaves more to wish for. *Internet protocol television* (IPTV) is being deployed on a wider extent and the transition to *high definition television* (HDTV) resolution is ongoing. In the longer run, 3D video and *free-viewpoint video* (FVV) services will also be offered. This development produces challenges for computer networks of all sizes, from small LANs to the whole Internet.

The employment of multicast transmission can reduce the resource demands of services where some content is simultaneously transmitted to a number of users. The reason is that the receivers of a multicast session share the re-

sources through a common transmission tree, where data are only transmitted once along each branch. Nevertheless, multicast transmission is not deployed to its full extent.

In [13], we therefore aimed at creating an incentive for the use of multicast transmission. The proposal was a general definition of how the bandwidth should be distributed fairly between competing multicast and unicast sessions. In short, the definition takes the number of receivers into consideration, which is beneficial for multicast sessions.

If the transmission costs for multicast sessions also were favorable when compared to those of unicast, this would create another incentive for the employment of multicast. In this paper, we therefore study how the transmission costs of multicast sessions should be allocated to achieve this goal. This work is an extension of that presented at the IARIA ICDDT 2007 conference [15] and in [14].

Henceforth, *costs* always refers to the costs associated with the actual transmission, i.e. costs for network resources such as links and routers, or in reality, the fees that the *Internet service providers* (ISPs) are charging. The cost of the delivered content is strictly excluded throughout this work.

To begin with, we adopt the fundamental assumption made by Herzog *et al.* in [8], that the cost of a multicast tree should be assigned to the receivers and not to the source. The reason is that multicast transmission is receiver initiated and that the service primarily is of use to the receivers, since the sources typically are streaming servers. The three basic requirements; no positive transfers, voluntary participation, and consumer sovereignty, are also sustained.

Further, we believe that fair cost allocation should be based on resource usage. This is likely to make the resource utilization more effective. With a flat-rate policy, there are no incentives for limiting the resource usage, as long as it is maintained within the postulated limit.

As an example, in everyday life, the expectation is that a train ticket will cost less than an air ticket. In addition, short domestic flights are expected to cost less than longer international flights. Furthermore, a shared cab is cheaper

per capita than a private one. The higher costs involved in more exclusive services together with a limited budget, probably accounts for the most common reason why people do not travel more, further, and faster, etc. A season ticket or the like, i.e. a flat rate policy, works against this incentive. Although, there might exist other motives, such as environmental awareness etc.

For data transmission over computer networks, the two major resource-related factors, which might differ between receivers, also relate to distance and quality. Namely the transmission path and the *quality of service* (QoS) requirements. As an example, choosing a server that is geographically close and settling for a low quality service would reduce the resource usage. This also holds for multicast receivers, but here the “shared-cab” aspect comes into play as well. Connecting to a multicast tree with many receivers in the vicinity will also save resources.

In Section 2 and 3 we describe existing cost-allocation mechanisms for multicast traffic. These mechanisms are then studied in Section 4, and the finding is that none take all of the aforementioned factors into consideration. A terminology for cost-allocation mechanisms that targets multi-rate multicast sessions is then introduced in Section 5, whereupon two new cost-allocation mechanisms are proposed in Section 6. The conclusions are presented in Section 7 together with some possible future research topics.

## 2. Existing cost-allocation mechanisms

In this section, a number of cost-allocation mechanisms for cost sharing among multicast receivers are outlined. These are a selection of existing mechanisms, other proposals for example include [5] and [3]. However, some of the terminology associated with cost sharing among multicast receivers is firstly introduced.

### 2.1. Terminology for multicast cost sharing

This section outlines the notations for cost sharing among multicast receivers, originally introduced in [8].

The number of receivers upstream and downstream respectively for a particular link are denoted by  $n_u$  and  $n_d$ . The receivers downstream of a link are those receivers whose transmission paths from the source traverse that link. The receivers upstream of a link are somewhat less intuitively defined as the receivers who are not located downstream of that link. In the multicast tree of Figure 1, where  $t$  is the transmitter, receivers  $r_1$ ,  $r_2$  and  $r_3$  are located downstream of link  $l$ , whereas receivers  $r_4$  through  $r_7$  are upstream of link  $l$ . The part of the cost of the link allocated to the upstream receivers is described by the function  $F_u(n_u, n_d)$ , whereas  $F_d(n_u, n_d)$  represents the part of the cost that is allocated to the downstream receivers.




Figure 1. A multicast transmission tree with seven receivers.

Multicast sessions that support multiple *quality of service* (QoS) levels are also covered in [8]. The shares of the total cost allocated to the upstream and downstream receivers requesting QoS level  $i$ , are denoted by  $F_u^i(z_u, z_d)$  and  $F_d^i(z_u, z_d)$  respectively. However, the terms  $z_u$  and  $z_d$  are not defined.

### 2.2. The edge-pricing paradigm

Pricing and cost allocation in computer networks are treated extensively by Shenker *et al.* in [12]. They initiate their discussion with pricing based on estimated congestion conditions. The reason being the high complexity associated with the computation of the actual prevailing congestion conditions and the consequence is basically QoS-sensitive time-of-day pricing. They then claim that differentiated pricing based on estimated congestion conditions can be exchanged for differentially priced QoS classes. When the estimated congestion probability is low, even cheaper QoS classes will perform well. Users can therefore adapt their costs by monitoring and changing QoS classes.

Shenker *et al.* further propose that the pricing, aside from the QoS class, only should depend on the locations of the source and destination. The costs of the actual transmission path are approximated using the costs of the expected path. Consequently, the prices are based upon the estimated congestion conditions along the expected transmission path from the source to the destination. If information about congestion conditions is gathered at the edges of the network of an ISP, it should be possible to determine the price of a session at the access point. For connections that traverse the borders between different ISPs, the ISPs must purchase the service from each other in the same manner that regular users purchase service. This solution is called the *edge-pricing paradigm*.

Multicast traffic causes a challenge for the edge-pricing

paradigm, because a multicast destination address is merely a logical name and does not identify the individual receivers of the multicast group. The only information about multicast sessions that is present in a router node is regarding the next hop(s). It is therefore impossible to estimate the multicast tree at the access points. Shenker *et al.* propose control messages to be sent when new receivers join a multicast group. These messages should be forwarded along the reverse multicast tree to the access point of the source, where the cost of the tree may be approximated. The ISPs would process the control messages at the edges of their network and thereby extract adequate information. An alternative solution is to record the cost of each link within the control messages.

Shenker *et al.* also have a general discussion relating to cost sharing among multicast receivers. However, they do not propose any cost-allocation mechanism.

### 2.3. Single QoS cost allocation

In [8], Herzog *et al.* present an extensive work regarding how the costs of multicast trees should be split among the receivers. They present a number of cost-allocation mechanisms, of which the *equal tree split* (ETS) and *equal link split downstream* (ELSD) mechanisms are given the most attention.

The ELSD cost-allocation mechanism splits the cost of each link in the tree evenly between the downstream receivers. Using the notations introduced in subsection 2.1, the part of the cost of the link allocated to the upstream receivers can be described as

$$F_u(n_u, n_d) = 0, \quad (1)$$

whereas the part of the cost allocated to each downstream receivers becomes

$$F_d(n_u, n_d) = \frac{1}{n_d}. \quad (2)$$

The ETS cost-allocation mechanism splits the cost of the entire transmission tree uniformly amongst all the receivers. Using the same notations, we obtain

$$F_u(n_u, n_d) = F_d(n_u, n_d) = \frac{1}{n_u + n_d}. \quad (3)$$

### 2.4. QoS-based cost allocation

If the transmitted data are hierarchically encoded and marked and the router nodes employ priority dropping, users may choose to subscribe to a service although they cannot utilize the entire data rate transmitted by the source. The most obvious reason behind such limitations are network connections with low capacity. When the transmitted content is real-time video, another limiting factor might

be the rendering capacity of the receiving device. In either case, these users do not utilize the entire bandwidth allocated to a multicast session, at least not on all of the links along their transmission path.

In [8], Herzog *et al.* observe that this should affect the cost allocation of multicast sessions, but they do not propose any specific cost-allocation mechanism for these scenarios. Using the terminology of subsection 2.1, they do however point out that if the cost-allocation functions fulfill the following condition,

$$\sum_{i=1}^I (z_u^i \cdot F_u^i(z_u, z_d) + z_d^i \cdot F_d^i(z_u, z_d)) = 1, \quad (4)$$

the costs associated with the link in question are fully allocated among the receivers.

Liu *et al.* study usage-based pricing and cost sharing of multicast traffic in [9]. They propose a cost-allocation mechanism, whose cost sharing they state “is proportional to individual members resource requirements, should a unicast service be used”. The receivers are divided into categories depending on their requested QoS level. The costs associated with a particular category are then aggregated over the entire multicast tree, but only split among receivers obtaining that QoS level or higher, in an ETS fashion. Henceforth, this cost-allocation mechanism is therefore referred to as *QoS-dependent ETS* (QoS-D ETS).

## 3. Game-theoretic cost-allocation mechanisms

Many researchers have considered the bandwidth-allocation and pricing process from a game-theoretic perspective. Somewhat simplified, this implies that potential users place bids which reflect what the service is worth to them. The ISP then allocates the resources according to these bids. Some basic notions of game theory that are introduced in [11] are outlined in 3.1, followed by two game-theoretic cost-allocation mechanisms. Other works on the same subject are [4] and [2].

### 3.1. Game-theoretic notions

A cost-allocation mechanism in which the costs allocated to the users exactly match the cost of the service, is called *budget balanced*. A user’s *welfare* can be described as the satisfaction after obtaining a service for a certain cost. An *efficient* cost-allocation mechanism chooses to serve the set of users that maximizes the aggregated welfare of all the users.

Assume that a user is part of a user set that is a subset of a larger set of users. Then a cost-allocation mechanism is *cross-monotonic* if for all such user sets, the cost allocated

to the user when the larger set is served, is lower or equal in comparison to when the smaller set is served.

It is reasonable to assume that users are selfish and place bids that maximize their probable welfare. A cost-sharing mechanism is *strategyproof* if users maximize their welfare by placing bids that truthfully correspond to how much the service is worth to them. *Group strategyproof* is a harder criterion that requires the cost-allocation mechanism to be resistant against groups of users who jointly place their bids in an attempt to increase their welfares.

Another contribution of [11], is the establishing of the following three basic requirements:

- *no positive transfers* – no user is paid to obtain a service
- *voluntary participation* – no user is forced to obtain a service
- *consumer sovereignty* – no user is refused a service if their bid is sufficiently high

According to [6], there are two cost-allocation mechanisms that are naturally strategyproof and adhere to these basic requirements, the *marginal-cost* (MC) and *Shapley-value* (SH) mechanisms. Further, it is stated that these are the two most appropriate mechanisms for cost sharing among multicast receivers.

### 3.2. The Shapley-value mechanism

The SH cost-allocation mechanism is the game-theoretical equivalent to ELSD. It splits the cost of a network link equally between all receivers that are located downstream [6]. The SH mechanism is group strategyproof and budget balanced. However, it is not efficient but has the smallest maximum loss of welfare among the budget-balanced mechanisms.

### 3.3. The marginal-cost mechanism

As described in [11], the MC mechanism essentially charges the marginal cost to the users, that is the cost of providing the service to all users minus the cost of providing the service to all but the user in question. It therefore has the characteristic that it treats equals equally, that is if two receivers give rise to the same marginal cost and place identical bids, they are allocated the same amount of resources and are charged the same cost. Further, the MC mechanism is efficient but not budget balanced nor group strategyproof.

In [1], the MC mechanism is applied to multicast sessions that support multiple rates. The *split session* and *layered* paradigms are studied, but only the layered paradigm is somewhat relevant here, since a split session basically implies separate transmissions of different QoS levels, i.e. the

problem associated with multiple QoS levels is divided into a number of problems, each with a single QoS level.

The layered paradigm, thoroughly described in [10], utilizes hierarchically encoded data, which is divided into QoS layers that are transmitted to individual multicast groups. The receivers consequently join multicast groups with QoS layers that can be combined into the desired QoS level. The layered paradigm therefore inherently implies that costs are separated according to QoS requirements.

### 3.4. Comparison of SH and MC mechanisms

In [7], both the SH and MC cost-allocation mechanisms are implemented and experiments are carried out. The MC is shown to generate a smaller revenue, which is not surprising since it is not budget balanced. On the other hand, the MC mechanism is faster than the SH mechanism.

In [6], it is observed that the MC mechanism only requires two messages per link in the multicast tree, whereas the number of messages required for the SH mechanism is of the order of the square of the number of links.

## 4. Evaluation of existing mechanisms

In this section, the cost-allocation mechanisms outlined in Section 2 and 3 are evaluated based on their attractiveness to the receivers. Important parameters are the magnitude of the costs and how fairly the costs are distributed.

### 4.1. The edge-pricing paradigm

The edge-pricing paradigm [12], briefly described in subsection 2.2, possesses some attractive properties, and it appears to be based upon sound approximations. However, the authors do not specify the pricing policy to be used. This decision is left to the individual ISPs. There are two main classes of pricing policies; usage-based policies where users are charged based on their actual usage, and capacity-based or flat-rate policies, where the users pay for the desired capacity. The choice, in this case, was to focus on usage-based pricing policies, since they are more favorable to multicast sessions and also might be considered to be fairer.

### 4.2. Single QoS cost allocation

For usage-based pricing policies, the cost of a multicast session should be divided among the receivers. The receivers in a multicast group have unique transmission paths per definition, otherwise they would have been positioned at the same location. As outlined in subsection 2.3, Herzog *et al.* propose a couple of cost-allocation mechanisms that are based upon the individual receivers' transmission paths [8]. However, there is a second factor that might affect

the amount of resources that are utilized by the individual receivers, namely the QoS requirements.

### 4.3. QoS-based cost allocation

As stated in subsection 2.4, users may choose to subscribe to a service although they cannot utilize the entire data rate transmitted by the source. These users do not use the entire bandwidth allocated to a multicast session, and should therefore, from a usage-based pricing perspective, be allocated a smaller share of the costs.

Although the work of Herzog *et al.* presented in [8] is extensive, the case involving individual receivers of a multicast group requesting different levels of QoS is covered on less than half a page. The discussion is very general and no specific cost-allocation mechanism is proposed for these scenarios.

The QoS-D ETS cost-allocation mechanism described by Liu *et al.* in [9] does however represent this approach. The costs corresponding to each QoS level are aggregated over the entire multicast tree, and divided uniformly among the receivers obtaining that level or higher. Thus, the lengths of the individual transmission paths are not taken into consideration. The statement in [9] concerning the cost sharing being proportional to the individual receivers' resource requirements, if unicast had been used, is therefore not strictly true.

### 4.4. Game-theoretic approaches

In game-theoretic approaches, the bandwidth allocation is incorporated with the pricing procedure. However, we aim for a cost-allocation mechanism that can fairly distribute the costs of any bandwidth allocation. The game-theoretic mechanisms are therefore ruled out.

### 4.5. Section summary

The game-theoretic approaches do not support cost-allocation of arbitrary bandwidth allocations, and none of the pure cost-allocation mechanisms takes both the transmission path and the QoS requirements into consideration. Hence, the mechanisms do not fully reflect the resource usage, and consequently there is room for improvements.

## 5. Terminology for multicast cost sharing

As mentioned in subsection 2.4, the notations for cost-allocation functions targeting multicast sessions with differentiated QoS levels, introduced by Herzog *et al.* in [8] and outlined in subsection 2.1, are not well defined. Thus, the decision was made to interpret and extend the terminology, in order to better suit multicast sessions that provide

multiple QoS levels. This will prove to be useful in the following section, where two new cost-allocation mechanisms are proposed.

We define  $n_u^q$  and  $n_d^q$  to be the number of upstream and downstream receivers of the  $q^{\text{th}}$  QoS level ( $QoS^q$ ), and let  $z_u^q$  and  $z_d^q$  denote the total number of upstream and downstream receivers utilizing the information corresponding to  $QoS^q$ . That is,

$$z_u^q = \sum_{x=q}^Q n_u^x$$

and

$$z_d^q = \sum_{x=q}^Q n_d^x,$$

given that there are  $Q$  available QoS levels. We also define the vectors

$$\mathbf{z}_u = \{z_u^1, z_u^2, \dots, z_u^Q\}$$

and

$$\mathbf{z}_d = \{z_d^1, z_d^2, \dots, z_d^Q\}.$$

Further, Herzog *et al.* not only allow the cost-allocation functions to control the division of the costs between receivers requesting the same QoS level, but also the distribution of the total cost among the different QoS levels. On the contrary, our opinion is that the cost-allocation functions should be general and not influence the distribution of the cost among the QoS levels. This distribution should instead fully reflect the resource requirements of each QoS level and the corresponding pricing made by the ISP in question.

Consequently, the cost vector

$$\mathbf{c} = \{c^1, c^2, \dots, c^Q\}$$

is introduced, where the additional costs for supporting  $QoS^q$  on a particular link during a specific period of time, when compared to those of  $QoS^{q-1}$ , are denoted by  $c^q$ . These costs should reasonably be split among the receivers requiring  $QoS^q$  or higher, and two cost-allocation subfunctions,  $f_u^q(z_u^q, z_d^q)$  and  $f_d^q(z_u^q, z_d^q)$ , are introduced for this purpose. These subfunctions describe the shares of the additional costs, for supporting  $QoS^q$  level, that should be allocated to the receivers of  $QoS^q$  or higher, both upstream and downstream of the link in question. The total cost that is to be allocated to the upstream and downstream receivers of  $QoS^q$  may now be written as

$$C_u^q(\mathbf{z}_u, \mathbf{z}_d, \mathbf{c}) = \sum_{x=1}^q f_u^x(z_u^x, z_d^x) c^x \quad (5)$$

and

$$C_d^q(\mathbf{z}_u, \mathbf{z}_d, \mathbf{c}) = \sum_{x=1}^q f_d^x(z_u^x, z_d^x) c^x, \quad (6)$$

respectively.

The two cost-allocation functions  $C_u^q(\mathbf{z}_u, \mathbf{z}_d, \mathbf{c})$  and  $C_d^q(\mathbf{z}_u, \mathbf{z}_d, \mathbf{c})$  represent the actual cost, whereas the original cost-allocation functions  $F_u(z_u, z_d)$  and  $F_d(z_u, z_d)$  described the fraction of the total cost to be allocated to the users. The condition (4), regarding full cost allocation, is therefore no longer valid. Instead, for the costs corresponding to each QoS level to be fully allocated, the following equation

$$z_u^q \cdot f_u^q(z_u^q, z_d^q) + z_d^q \cdot f_d^q(z_u^q, z_d^q) \geq 1, \quad (7)$$

must be fulfilled for all integers  $q$  between one and  $Q$ , where  $Q$  is the highest QoS level with a receiver downstream of the link in question.

If equation (7) is an equality for all integers  $q$  between one and  $Q$ , this guarantees that the sum of all allocated costs equals the sum of the costs according to equation (8), which means that the cost-allocation mechanism is budget balanced.

As an example, consider the QoS-D ETS cost-allocation mechanism described in subsection 2.4. Using the terminology introduced in this section, it is represented by cost-allocation subfunctions corresponding to the cost-allocation functions of the ETS mechanism (3)

$$f_u^q(z_u^q, z_d^q) = f_d^q(z_u^q, z_d^q) = \frac{1}{z_u^q + z_d^q}.$$

Consequently

$$z_u^q \cdot f_u^q(z_u^q, z_d^q) + z_d^q \cdot f_d^q(z_u^q, z_d^q) = z_u^q \frac{1}{z_u^q + z_d^q} + z_d^q \frac{1}{z_u^q + z_d^q} = \frac{z_u^q + z_d^q}{z_u^q + z_d^q} = 1,$$

and the QoS-D ETS mechanism is therefore budget balanced according to equation (8).

## 6. Fair cost-allocation strategies

The evaluation of existing cost-allocation mechanisms in Section 4 was concluded with the realization that none of them were satisfactorily fair. The reason was that, at most, they consider one of the two main factors affecting the resource usage, i.e. the transmission path and the QoS requirements. Using the terminology introduced in Section 5, a new cost-allocation mechanism, which takes both these factors into consideration, is proposed in subsection 6.1.

Although the aim of this mechanism is to achieve optimum fairness, it might have one, possibly severe, shortcoming: Optimum fairness may not be the primary interest of the receivers, if it occurs at the expense of higher costs. If poor and greedy receivers get a discount on the service, it may actually become cheaper for the rest of the receivers. An alternative mechanism is therefore proposed in subsection 6.2.

### 6.1. QoS-differentiated link split downstream

The first proposal is designed to perform perfectly fair cost allocations, taking into consideration both the transmission path and the QoS requirements. It builds on the ELSD cost-allocation mechanism, presented by Herzog *et al.* in [8], but is enhanced to support differentiated QoS levels.

The cost-allocation subfunctions therefore correspond to equations (1) and (2), and become

$$f_u^q(z_u^q, z_d^q) = 0 \quad (9)$$

and

$$f_d^q(z_u^q, z_d^q) = \frac{1}{z_d^q}, \quad (10)$$

respectively. This gives that

$$z_u^q \cdot f_u^q(z_u^q, z_d^q) + z_d^q \cdot f_d^q(z_u^q, z_d^q) = z_u^q \cdot 0 + z_d^q \frac{1}{z_d^q} = \frac{z_d^q}{z_d^q} = 1,$$

and the cost-allocation mechanism is consequently budget balanced according to equation (8).

Substituting equations (9) and (10) into (5) and (6), the main cost-allocation functions for receivers of  $QoS^q$  become

$$C_u^q(\mathbf{z}_u, \mathbf{z}_d, \mathbf{c}) = 0 \quad (11)$$

and

$$C_d^q(\mathbf{z}_u, \mathbf{z}_d, \mathbf{c}) = \sum_{x=1}^q \frac{c^x}{z_d^x}. \quad (12)$$

We call the cost-allocation mechanism described by equations (11) and (12), the *QoS-differentiated link split downstream* (QoS-D LSD) mechanism.

#### 6.1.1. Bandwidth-differentiated link split downstream.

As observed in [8], in the extreme case, each receiver will have a QoS level of its own. This can be taken one step further, by assuming the bandwidth to be the predominant cost factor and considering the bandwidth consumption as a direct function of the QoS level. Let us also assume that the bandwidth is uniformly priced and costs  $c$  monetary units (MU) per *bitrate unit* (BU) and *time unit* (TU).

Let  $\mathbf{b}$  be a vector whose first element  $\mathbf{b}[0]$  is 0 and the  $n_d$  following elements are the receiving rates of the receivers downstream of the link in question, sorted in ascending order. The total cost per TU, allocated to the downstream receiver obtaining the  $q^{\text{th}}$  smallest bandwidth, may now be rewritten as

$$C_d^q(n_d, \mathbf{b}) = c \sum_{x=1}^q \frac{\mathbf{b}[x] - \mathbf{b}[x-1]}{n_d - x + 1}. \quad (13)$$

$$\begin{aligned}
& \sum_{q=1}^Q (n_u^q \cdot C_u^q(\mathbf{z}_u, \mathbf{z}_d, \mathbf{c}) + n_d^q \cdot C_d^q(\mathbf{z}_u, \mathbf{z}_d, \mathbf{c})) \\
&= \sum_{q=1}^Q \left( n_u^q \cdot \sum_{x=1}^q f_u^x(z_u^x, z_d^x) c^x + n_d^q \cdot \sum_{x=1}^q f_d^x(z_u^x, z_d^x) c^x \right) \\
&= \left( n_u^1 \cdot f_u^1(z_u^1, z_d^1) c^1 + n_d^1 \cdot f_d^1(z_u^1, z_d^1) c^1 \right) \\
&+ \left( n_u^2 \cdot (f_u^1(z_u^1, z_d^1) c^1 + f_u^2(z_u^2, z_d^2) c^2) + n_d^2 \cdot (f_d^1(z_u^1, z_d^1) c^1 + f_d^2(z_u^2, z_d^2) c^2) \right) + \dots \\
&\dots + \left( n_u^Q \cdot (f_u^1(z_u^1, z_d^1) c^1 + f_u^2(z_u^2, z_d^2) c^2 + \dots + f_u^Q(z_u^Q, z_d^Q) c^Q) \right. \\
&\quad \left. + n_d^Q \cdot (f_d^1(z_u^1, z_d^1) c^1 + f_d^2(z_u^2, z_d^2) c^2 + \dots + f_d^Q(z_u^Q, z_d^Q) c^Q) \right) \tag{8} \\
&= \left( f_u^1(z_u^1, z_d^1) c^1 \cdot (n_u^1 + n_u^2 + \dots + n_u^Q) + f_d^1(z_u^1, z_d^1) c^1 \cdot (n_d^1 + n_d^2 + \dots + n_d^Q) \right) \\
&+ \left( f_u^2(z_u^2, z_d^2) c^2 \cdot (n_u^2 + n_u^3 + \dots + n_u^Q) + f_d^2(z_u^2, z_d^2) c^2 \cdot (n_d^2 + n_d^3 + \dots + n_d^Q) \right) + \dots \\
&\quad \dots + \left( f_u^Q(z_u^Q, z_d^Q) c^Q \cdot n_u^Q + f_d^Q(z_u^Q, z_d^Q) c^Q \cdot n_d^Q \right) \\
&= \sum_{q=1}^Q \left( f_u^q(z_u^q, z_d^q) c^q \cdot \sum_{x=q}^Q n_u^x + f_d^q(z_u^q, z_d^q) c^q \cdot \sum_{x=q}^Q n_d^x \right) \\
&= \sum_{q=1}^Q c^q \cdot (z_u^q \cdot f_u^q(z_u^q, z_d^q) + z_d^q \cdot f_d^q(z_u^q, z_d^q)) = \sum_{q=1}^Q c^q
\end{aligned}$$

The *bandwidth-differentiated link split downstream cost allocation* performed by equation (13) is only a special case of the QoS-D LSD mechanism.

**6.1.2. A cost-allocation example.** As a small example of the QoS-D LSD mechanism, let us study how equation (13) allocates the cost of link  $l$  in Figure 1, where  $t$  is the transmitter and  $r_1$  through  $r_7$  are the receivers. For simplicity, we assume that receiver  $r_i$  obtains  $i$  BU for one TU, and that the bandwidth on link  $l$  costs one MU per BU and TU. Now we have

$$\begin{aligned}
c &= 1 \\
n_d &= 3 \\
\mathbf{b} &= \{0, 1, 2, 3\},
\end{aligned}$$

which when substituted into equation (13) give the cost of link  $l$  being allocated to receiver  $r_1$ ,  $r_2$ , and  $r_3$  as follows,

$$C_d^1(n_d, \mathbf{b}) = \sum_{x=1}^1 \frac{\mathbf{b}[x] - \mathbf{b}[x-1]}{4-x} = \frac{1}{3} \text{ MU},$$

$$C_d^2(n_d, \mathbf{b}) = \sum_{x=1}^2 \frac{\mathbf{b}[x] - \mathbf{b}[x-1]}{4-x} = \frac{1}{3} + \frac{1}{2} = \frac{5}{6} \text{ MU},$$

and

$$\begin{aligned}
C_d^3(n_d, \mathbf{b}) &= \sum_{x=1}^3 \frac{\mathbf{b}[x] - \mathbf{b}[x-1]}{4-x} = \frac{1}{3} + \frac{1}{2} + \frac{1}{1} \\
&= \frac{11}{6} \text{ MU}.
\end{aligned}$$

If we, similarly, calculate the total costs allocated to receiver  $r_1$ ,  $r_2$ , and  $r_3$ , link by link from the source, they be-

come

$$\left(\frac{1}{4}\right) + \left(\frac{1}{3}\right) + \left(\frac{1}{1}\right) = \frac{19}{12} \text{ MU},$$

$$\left(\frac{1}{4} + \frac{1}{3}\right) + \left(\frac{1}{3} + \frac{1}{2}\right) + \left(\frac{2}{2}\right) + \left(\frac{2}{1}\right) = \frac{53}{12} \text{ MU},$$

and

$$\begin{aligned}
&\left(\frac{1}{4} + \frac{1}{3} + \frac{1}{2}\right) + \left(\frac{1}{3} + \frac{1}{2} + \frac{1}{1}\right) + \left(\frac{2}{2} + \frac{1}{1}\right) + \left(\frac{3}{1}\right) \\
&= \frac{95}{12} \text{ MU},
\end{aligned}$$

respectively. To make the calculations easier to follow, the costs are presented for every bandwidth interval, and costs arising from the same link are grouped together by parentheses.

The costs allocated to all the seven receivers in the multicast tree are presented in Table 1, together with the corresponding costs produced by the ETS, ELSD, and QoS-D ETS cost-allocation mechanisms.

The ETS and ELSD mechanisms were not designed with differentiated QoS demands in mind. Both these mechanisms will therefore generally allocate disproportionately large parts of the cost to receivers with low QoS demands. The ETS mechanism simply splits the aggregated cost of the entire multicast tree equally among all the receivers, and is therefore also unfair towards receivers with short transmission paths. The ELSD mechanism only splits the link costs among downstream receivers, and the receivers that are treated most unfairly are consequently those with low QoS demands, compared to the receivers with whom they share the links. Examples of such mistreated receivers are consequently  $r_1$ ,  $r_2$ , and  $r_5$ .

**Table 1. The obtained bitrates in BUs of the seven receivers in the example, together with the costs in MUs, allocated by the ETS, ELSD, QoS-D ETS, and QoS-D LSD cost-allocation mechanisms.**

receiver	rate	ETS	ELSD	QoS-D ETS	QoS-D LSD
$r_1$	1	7.29	3.00	1.71	1.58
$r_2$	2	7.29	5.50	3.55	4.42
$r_3$	3	7.29	6.50	5.55	7.92
$r_4$	4	7.29	5.00	7.30	6.08
$r_5$	5	7.29	10.3	8.96	9.17
$r_6$	6	7.29	11.3	11.0	11.7
$r_7$	7	7.29	9.33	13.0	10.2

The QoS-D ETS mechanism performs differently, as it is now the receivers with short transmission paths, such as  $r_4$  and  $r_7$ , that are treated unfairly. The situation is worst for  $r_7$ , which obtains the highest QoS level, and therefore has to share the costs of the entire multicast tree.

## 6.2. Bid-based link split downstream

As mentioned previously, the proposed QoS-D LSD cost-allocation mechanism attempts to achieve optimum fairness, but it has one possibly severe shortcoming: Optimum fairness may not be the primary interest of the receivers if it is at the expense of higher costs. If poor and greedy receivers get a discount on the service, it may actually become cheaper for the rest of the receivers. Here we further investigate this issue and propose an alternative cost-allocation mechanism that solves the shortcoming.

We start by drawing a parallel to an everyday situation. Children and/or retired people often receive a discount on the entrance fee to sport events, festivals, and museums etc. Most people are willing to accept this since it typically does not negatively affect their fees. As long as the events are not sold out, the economy of the organizers might actually benefit from this, and thereby allow them to also lower the standard fees<sup>1</sup>.

However, if the scenario was the opposite and the attendance of discounted groups had a negative influence on standard fees, i.e. forcing the regular visitors to subsidize those on discounted rates, few would be happy about accepting such a system. Consumer goods are seldom discounted in this manner, since they are associated with specific material and production costs.

<sup>1</sup>If any organizers actually do this in reality is a completely different question.

**Table 2. Possible outcomes of a placed bid, with a certain maximum cost, for the BB LSD cost-allocation mechanism.**

relative size of the maximum cost	served	allocated cost
$\max \text{ cost} < \text{additional cost}$	no	–
$\text{additional cost} \leq \max \text{ cost} < \text{fair share}$	yes	max cost
$\text{fair share} \leq \max \text{ cost}$	yes	fair share

If we look at the game-theoretic approaches of Section 3, the SV mechanism allocates the costs in a LSD manner, and therefore shares the aforementioned shortcoming. The MC mechanism on the other hand does not require the receivers to cover more than their marginal cost. It is consequently not budget balanced, and may thereby produce a financial deficit for the ISPs.

We propose a bid-based cost-allocation mechanism, where fair cost allocation according to the QoS-D LSD mechanism is retained as the target. However, bids that do not cover the receivers' fair shares of the costs, but do cover at least the *additional cost* associated with receivers' requests, are also accepted. That is, the additional cost for providing the receiver with the requested service, compared to the cost of providing the service to the existing set of receivers.

The main difference between marginal cost and additional cost is that the latter is dependent upon the order of the arrival of the bids which, in turn, guarantees that the proposed mechanism is budget balanced. However, although an expansion of the user set never causes increased costs for users within the original set, the mechanism is not cross monotonic, since it only applies to ordered sets of users.

A placed bid consequently leads to one of the outcomes described in Table 2. The fair share is the cost calculated according to the principles of the QoS-D LSD mechanism, with the addendum that if some poor receivers are discounted, these costs have to be carried by the wealthier receivers. The costs not covered by a receiver are distributed between the affected links and QoS levels of the existing transmission tree, proportional to that receiver's fair cost shares, and are split among the higher-bidding receivers utilizing these resources. The proposed mechanism is called *bid-based link split downstream* (BB LSD).

**6.2.1. Bid structure.** There are a number of mandatory parameters that a bid must contain to make the BB LSD possible, namely:

- the maximum acceptable cost of the transmission
- the requested duration of the transmission



**Table 3. The bids of the receivers in the example in subsection 6.3. The maximum cost is measured in MU.**

receiver	requested QoS	maximum cost
$r_1$	$QoS^1$	25
$r_2$	$QoS^2$	25
$r_3$	$QoS^3$	80
$r_4$	$QoS^4$	100


- the requested QoS level of the transmission
- the *time to live* (TTL) of the bid

It is insufficient to replace the maximum cost and requested duration with a maximum cost per TU. This would prevent the calculation of other receivers' maximum costs, since these are affected by receivers who leave the service prematurely. There is also a possibility of non-recurrent costs associated with setting up the service. The bid TTL is required since most users are only interested in a particular service if it can be started within a given amount of time.

A receiver may request a service at a particular price, but be willing to settle for a poorer QoS level at a lower price if the main bid cannot be accepted. The main bid could then possibly remain effective during its TTL, in case the costs associated with it were to be reduced. We observe that there may be as many subbids as there are QoS levels, but do not discuss these composite bids any further.

**6.2.2. Strategyproofness.** The BB LSD mechanism is not strategy proof. There is an obvious risk that users place dishonestly low bids, i.e. bids that do not correspond to their estimated value of the service, in an attempt to find the minimum cost of the service. To avoid this destructive behavior for the system, we propose an upper limit on the bid frequency of any particular receiver. This might not make the mechanism strategy proof, but it should make users more honest, since a lower bid equals a higher risk of missing out on the service for a particular amount of time.

The problem of finding a sufficient maximum bid frequency is a weighing of the honesty of the bids against the adaptability of the mechanism. It is possible that the economic prerequisites of a receiver change for the better after a low bid has been placed. An alternative to a fixed maximum bid frequency, is to exponentially increase the period of time until a new bid might be placed or considered.



**Figure 2. The multicast transmission tree of the example in subsection 6.3.**

### 6.3. A cost-allocation example

The transmission tree in Figure 2 is used as an example in order to shed some light on the possible advantages of the BB LSD cost-allocation mechanism. The requested QoS levels are outlined in Table 3, together with the maximum total cost that the receivers are willing to pay for the service. For simplicity, assume that all requests concern the same duration, say 10 TU, and that the bandwidth on all links cost one MU per BU and TU. Further assume that the bitrate is the predominant cost factor and that  $QoS^q$  constantly requires  $q$  BU. The incremental cost of transmitting  $QoS^q$ , when compared to that of  $QoS^{q-1}$ , is consequently one MU/TU per link.

In the two first subsections, the QoS-D LSD and MC cost-allocation mechanisms are utilized to allocate the bandwidth and costs, and in the third subsection, these parameters are calculated according to the proposed BB LSD mechanism. For the latter mechanism, the order of arrival of the bids is essential. For simplicity, we base the order on the receiver numbers, and assume the arrivals of the bids to be sufficiently closely spaced in time for the requested transmissions to be considered simultaneously from a cost-sharing perspective. The results of the cost-allocation mechanisms are compared in the last subsection.

**6.3.1. Allocation according to QoS-D LSD.** We start by studying how the QoS-D LSD mechanism would allocate the cost of link  $l_2$ , under the assumption that all receivers are able to obtain the requested service at prices not exceeding their maximum costs. According to equation (12), receiver  $r_1$  will be charged

$$\frac{10}{3} \approx 3.33 \text{ MU}$$

for receiving  $QoS^1$ , since there are three receivers utilizing

this information. In the same manner, the cost of link  $l_2$  allocated to receivers  $r_2$  and  $r_3$ , which are requesting  $QoS^2$  respectively  $QoS^3$ , become

$$\frac{10}{3} + \frac{10}{2} \approx 8.33 \text{ MU}$$

and

$$\frac{10}{3} + \frac{10}{2} + \frac{10}{1} \approx 18.33 \text{ MU.}$$

The cost of each receiver can be calculated link by link from the source. The total costs of receivers  $r_1$  through  $r_4$  then become

$$\begin{aligned} & \left(\frac{10}{4}\right) + \left(\frac{10}{3}\right) + \left(\frac{10}{1}\right) \approx 15.83 \text{ MU,} \\ & \left(\frac{10}{4} + \frac{10}{3}\right) + \left(\frac{10}{3} + \frac{10}{2}\right) + \left(\frac{10}{2} + \frac{10}{1}\right) + \\ & \quad \left(\frac{10}{1} + \frac{10}{1}\right) \approx 44.17 \text{ MU,} \end{aligned} \quad (14)$$

$$\begin{aligned} & \left(\frac{10}{4} + \frac{10}{3} + \frac{10}{2}\right) + \left(\frac{10}{3} + \frac{10}{2} + \frac{10}{1}\right) + \\ & \left(\frac{10}{2} + \frac{10}{2} + \frac{10}{1}\right) + \left(\frac{10}{1} + \frac{10}{1} + \frac{10}{1}\right) \approx 79.17 \text{ MU,} \end{aligned} \quad (15)$$

and

$$\begin{aligned} & \left(\frac{10}{4} + \frac{10}{3} + \frac{10}{2} + \frac{10}{1}\right) + \left(\frac{10}{1} + \frac{10}{1} + \frac{10}{1} + \frac{10}{1}\right) \\ & \approx 60.83 \text{ MU,} \end{aligned}$$

respectively. To make the calculations easier to follow, the costs arising from the same link are grouped by parentheses.

Apparently, the assumption that all receivers are able to obtain the service, at a cost not exceeding their maximum limits, was false. Receiver  $r_2$  is only willing to pay 25 MU, but would be charged over 44 MU. It will therefore not obtain the service, and the rest of the receivers will consequently have to cover a larger part of the costs on the shared links. Receivers  $r_1$ ,  $r_3$ , and  $r_4$  will now be charged

$$\left(\frac{10}{3}\right) + \left(\frac{10}{2}\right) + \left(\frac{10}{1}\right) \approx 18.33 \text{ MU,}$$

$$\begin{aligned} & \left(\frac{10}{3} + \frac{10}{2} + \frac{10}{2}\right) + \left(\frac{10}{2} + \frac{10}{1} + \frac{10}{1}\right) + \\ & \left(\frac{10}{1} + \frac{10}{1} + \frac{10}{1}\right) + \left(\frac{10}{1} + \frac{10}{1} + \frac{10}{1}\right) \approx 103.33 \text{ MU,} \end{aligned}$$

respectively

$$\begin{aligned} & \left(\frac{10}{3} + \frac{10}{2} + \frac{10}{2} + \frac{10}{1}\right) + \left(\frac{10}{1} + \frac{10}{1} + \frac{10}{1} + \frac{10}{1}\right) \\ & \approx 63.33 \text{ MU.} \end{aligned}$$

Hence, the cost allocated to receiver  $r_3$  exceeds its bid of 80 MU, and it will also fail to obtain the requested service. The costs of receivers  $r_1$  and  $r_4$  are increased accordingly to

$$\left(\frac{10}{2}\right) + \left(\frac{10}{1}\right) + \left(\frac{10}{1}\right) = 25.00 \text{ MU}$$

and

$$\begin{aligned} & \left(\frac{10}{2} + \frac{10}{1} + \frac{10}{1} + \frac{10}{1}\right) + \left(\frac{10}{1} + \frac{10}{1} + \frac{10}{1} + \frac{10}{1}\right) \\ & = 75.00 \text{ MU,} \end{aligned}$$

respectively. Finally, all the costs are covered by the receivers' bids.

**6.3.2. Allocation according to MC.** The MC cost-allocation mechanism has received its name because it allocates the marginal cost to each user. The marginal cost of a user is the additional cost of providing the service to that user, when compared to the cost of providing the service to the remaining set of users.

In this example the marginal cost of receiver  $r_1$  corresponds to that of  $QoS^1$  on link  $l_3$ , i.e. 10 MU, since  $r_2$  and  $r_3$  also utilize  $QoS^1$  on the rest of the transmission path from the source to  $r_1$ . On link  $l_1$ ,  $QoS^1$  is also utilized by receiver  $r_4$ .

In the same manner, the marginal cost of receiver  $r_2$  is derived from the provision of  $QoS^2$  on link  $l_5$ , that is 20 MU. On the rest of the transmission path from the source to  $r_2$ ,  $QoS^2$  is shared by receiver  $r_3$ .

Receiver  $r_3$  is allocated the total cost for  $QoS^3$  on its last hop link  $l_6$ , which corresponds to 30 MU. Further, on links  $l_2$  and  $l_4$ ,  $r_3$  is the only receiver that utilizes  $QoS^3$ . It therefore has to cover the additional cost of  $QoS^3$ , when compared to that of  $QoS^2$ , on these links. This implies a cost of 10 MU per link. However,  $r_3$  does not have to contribute to the costs of  $l_1$ , since  $QoS^3$  is shared with receiver  $r_4$  on that link. The aggregated cost allocated to receiver  $r_3$  is consequently 50 MU.

Finally, receiver  $r_4$  is charged with the total cost of  $QoS^4$  on link  $l_7$  and the additional cost of  $QoS^4$  on link  $l_1$ . This adds up to a total of 50 MU, and all receivers will therefore be served since the maximum costs of their bids cover the allocated costs.

**6.3.3. Allocation according to BB LSD.** Now the proposed BB LSD cost-allocation mechanism is applied to the same example.

When the bid of receiver  $r_1$  is placed, its maximum cost of 25 MU is insufficient to cover the cost of the requested  $QoS^1$ , which is calculated to 30 MU over the three-link transmission path from the source. The bid is therefore not accepted, but remains effective, pending other bids that may share the costs.

When the bid of receiver  $r_2$  arrives, the costs associated with its request for  $QoS^2$  is 80 MU. The bid is on 25 MU, and can therefore not be accepted either, not even when considered jointly with the bid of  $r_1$ .

Then the bid of receiver  $r_3$  is placed. It concerns  $QoS^3$  and is worth 80 MU, whereas the cost for offering the service is 120 MU. The total cost for serving  $r_1$ ,  $r_2$ , and  $r_3$  would be 150 MU, whereas their joint means are calculated as being 130 MU. Separately considering  $r_1$  and  $r_3$ , or  $r_2$  and  $r_3$ , does not make the situation more favorable.

Finally, the bid of receiver  $r_4$  is placed. The costs for

the requested transmission to  $r_4$  is 80 MU, and the bid on 100 MU can therefore be accepted on its own. However, to decide what costs will actually be allocated to  $r_4$ , the bids of the other receivers must first be reconsidered.

Let us start by considering receiver  $r_2$ . The costs of the resources that  $r_2$  must cover in total, i.e. those of link  $l_5$ , are 20 MU according to the last parenthesis of equation (14). It therefore has 5 MU left to contribute to the cost sharing on the upstream links. These 5 MU will be split uniformly according to  $r_2$ 's fair shares of the costs on these links, which corresponds to the remaining first three parenthesis of (14). This results in

$$5 \cdot \frac{\frac{10}{4}}{\left(\frac{10}{4} + \frac{10}{3}\right) + \left(\frac{10}{3} + \frac{10}{2}\right) + \left(\frac{10}{2} + \frac{10}{2}\right)} \approx 0.52 \text{ MU}$$

for  $QoS^1$  on  $l_1$ , and in the same manner approximately 0.69 MU for  $QoS^2$  on  $l_1$  and  $QoS^1$  on  $l_2$ , and 1.03 MU for  $QoS^2$  on  $l_2$ , and  $QoS^1$  and  $QoS^2$  on  $l_4$ .

Receiver  $r_3$  must cover the entire 30 MU for link  $l_6$  and the remaining costs on  $l_4$ . Further, it also has to cover the additional cost of  $QoS^3$  on  $l_2$  together with the remaining cost for  $QoS^2$ . Consequently, there are approximately

$$80 - 30 - (30 - 2 \cdot 1.03) - (20 - 1.03) \approx 3.09 \text{ MU}$$

left on the bid of  $r_3$ . Split uniformly according to  $r_3$ 's remaining costs shares, which can be found in equation (15), this yields

$$3.09 \cdot \frac{\frac{10}{4}}{\left(\frac{10}{4} + \frac{10}{3} + \frac{10}{2}\right) + \left(\frac{10}{3}\right)} \approx 0.55 \text{ MU}$$

for  $QoS^1$  on link  $l_1$ , and in the same manner approximately 0.73 MU for  $QoS^2$  on  $l_1$ , 1.09 MU for  $QoS^3$  on  $l_1$ , and 0.73 MU for  $QoS^1$  on  $l_2$ .

Consequently, receiver  $r_1$  that only requested  $QoS^1$ , has to cover 10 MU on link  $l_3$  and the remaining costs on  $l_2$ , which is approximately

$$10 - 0.69 - 0.73 = 8.58 \text{ MU.}$$

On link  $l_1$ ,  $r_1$  will be charged with its own fair share of the costs, plus its share of the costs for  $QoS^1$  that are not covered by  $r_2$  and  $r_3$ . This adds up to

$$\frac{10}{4} + \frac{\frac{10}{4} - 0.52}{2} + \frac{\frac{10}{4} - 0.55}{2} \approx 4.47 \text{ MU.}$$

The total cost allocated to  $r_1$  thereby aggregates into approximately

$$10 + 8.58 + 4.47 = 23.05 \text{ MU.}$$

The remaining costs, which are allocated to receiver  $r_4$ , are calculated as being 40 MU for link  $l_7$ , and approximately

$$(10 - 4.47 - 0.52 - 0.52) + (10 - 0.69 - 0.73) + (10 - 1.09) + 10 = 31.98 \text{ MU}$$

**Table 4. The outcomes for the receivers with the QoS-D LSD, MC and BB LSD cost-allocation mechanisms. The costs are measured in MU.**

receiver	QoS-D LSD		MC		BB LSD	
	served	cost	served	cost	served	cost
$r_1$	yes	25.0	yes	10.0	yes	23.0
$r_2$	no	–	yes	20.0	yes	25.0
$r_3$	no	–	yes	50.0	yes	80.0
$r_4$	yes	75.0	yes	50.0	yes	72.0

**Table 5. The announced costs of the provided services and the generated incomes, both measured in MU, with the QoS-D LSD, MC and BB LSD cost-allocation mechanisms.**

	QoS-D LSD	MC	BB LSD
announced service costs	100	200	200
generated incomes	100	130	200

for link  $l_1$ , where each QoS level is accounted for separately. This gives a total cost for receiver  $r_4$  of approximately 71.98 MU.

**6.3.4. Comparison of results.** In Table 4, the outcomes for the receivers with the proposed BB LSD cost-allocation mechanism are presented together with them of MC and QoS-D LSD.

The most obvious difference between the BB LSD and QoS-D LSD mechanisms is that receivers  $r_2$  and  $r_3$  are served by BB LSD but not by QoS-D LSD, since they cannot fully cover their fair shares of the costs. As a consequence, the costs allocated to receivers  $r_1$  and  $r_4$  are somewhat lower for the BB LSD mechanism, where receiver  $r_3$  contributes to the cost sharing on links  $l_1$  and  $l_2$ . Another, more significant effect, which is apparent in Table 5, is that the income of the ISP is doubled through the use of the BB LSD mechanism.

The BB LSD and MC mechanisms serve the same user sets. However, all the receivers are allocated lower costs by using the MC mechanism, since it only charges the marginal costs. As can be seen in Table 5, the result is, if not a financial deficit, at least a 70 MU reduction of the ISP's revenue, when compared to the budget-balanced BB LSD mechanism.

## 7. Conclusion

This paper has aimed at more efficient usage of bandwidth in IP networks. The area that has been targeted is the slow deployment of multicast transmission. The proposal was to reduce the costs for users of multicast sessions. The cost reduction is brought about by the resource savings offered by the bandwidth sharing.

Fair cost sharing among multicast receivers has been addressed. This would favor the multicast receivers under the assumption that fair cost sharing should be based upon resource usage. Two major resource-related factors were observed; the transmission path and the bandwidth or QoS requirements. Existing cost-allocation mechanisms for multicast were evaluated, but none took both these parameters into consideration. The QoS-D LSD cost-allocation mechanism was therefore proposed. It considers both the transmission path and the QoS requirements, in order to achieve optimum fairness.

However, optimum fairness might not be in the best interest of the users, when it is at the expense of higher costs. An alternative cost-allocation mechanism, BB LSD, was therefore proposed. The BB LSD mechanism enables the users to place bids for a requested service, revealing their maximum acceptable cost. A bid that does not cover the user's fair share of the costs for the requested service is nevertheless accepted if it does cover at least the additional cost associated with the request. This guarantees that the BB LSD mechanism is budget balanced. The result is not only a possible reduction in the costs for the rest of the users, but also an increase in revenue for the ISPs, which are able to serve more users.

Unfortunately, the BB LSD mechanism is not strategy proof. To avoid users seeking the minimum cost by placing dishonestly low bids, an upper limit on the bid frequency of any particular receiver was therefore proposed. Another alternative would be an exponentially growing time out in the case of a rejected bid. This should make the users more honest, i.e. to bid closer to what the service is worth to them, since a lower bid equals a higher risk of missing out on the service.

### 7.1. Future work

Future research about cost-allocation mechanisms may involve the problem of finding a sufficient maximum bid frequency, or other procedures to mitigate the fact that the BB LSD mechanism is not strategy proof. Another alternative might be the search for a completely new mechanism that is naturally strategy proof and still possesses as many of the BB LSD mechanism's attractive properties as possible.

Further research topics are the implementation of the QoS-D LSD and BB LSD cost-allocation mechanisms, and

the process of actually charging the receivers with the allocated costs.

## Acknowledgment

The work was financed in part by the Regional Fund of the European Union and the County Administrative Board of Västernorrland.

## References

- [1] M. Adler and D. Rubenstein. Pricing multicasting in more flexible network models. *ACM Transactions on Algorithms*, 1(1):48–73, July 2005.
- [2] M. Bläser. Approximate budget balanced mechanisms with low communication costs for the multicast cost-sharing problem. In *Proceedings of 15<sup>th</sup> ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 625–195, New Orleans, LA, USA, January 2004.
- [3] A. Bueno, P. Vila, and R. Fabregat. Multicast extension of unicast charging for QoS services. In *Proceedings of 4<sup>th</sup> IEEE European Conference on Universal Multiservice Networks (ECUMN)*, pages 119–126, Toulouse, France, February 2007.
- [4] S. Chawla, D. Kitchin, U. Rajan, R. Ravi, and A. Sinha. Profit guaranteeing mechanisms for multicast networks. In *Proceedings of 4<sup>th</sup> ACM Conference on Electronic Commerce (EC)*, pages 190–191, San Diego, CA, USA, June 2003.
- [5] H. J. Einsiedler, P. Hurley, B. Stiller, and T. Braun. Charging multicast communications based on a tree metric. In *Proceedings of GI Multicast Workshop*, Braunschweig, Germany, May 1999.
- [6] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Elsevier Journal of Computer and System Sciences*, 63(1):21–41, August 2001.
- [7] N. Garg and D. Grosu. Performance evaluation of multicast cost sharing mechanisms. In *Proceedings of 21<sup>st</sup> IEEE International Conference on Advanced Networking and Applications (AINA)*, pages 901–908, Niagara Falls, Canada, May 2007.
- [8] S. Herzog, S. Shenker, and D. Estrin. Sharing the “cost” of multicast trees: An axiomatic analysis. *IEEE/ACM Transactions on Networking*, 5(6):847–860, December 1997.
- [9] C.-C. Liu, S.-C. Chang, and H.-H. Cheng. Pricing and fee sharing for point to multipoint and quality guaranteed multicast services. In *Proceedings of 7<sup>th</sup> IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, pages 255–260, Iwate, Japan, July 2000.
- [10] S. McCanne, V. Jacobson, and M. Vetterli. Receiver-driven layered multicast. *ACM SIGCOMM Computer Communication Review*, 26(4):117–130, October 1996.
- [11] H. Moulin and S. Shenker. Strategyproof sharing of sub-modular access costs: Budget balance versus efficiency. *Springer Journal on Economic Theory*, 18(3):511–533, 2001.

- [12] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *ACM SIGCOMM Computer Communication Review*, 26(2):19–43, April 1996.
- [13] P. Österberg and T. Zhang. Revised definition of multicast-favorable max-min fairness. In *Proceedings of 3<sup>rd</sup> IASTED International Conference on Communications and Computer Networks (CCN)*, pages 63–68, Lima, Peru, October 2006.
- [14] P. Österberg and T. Zhang. Bid-based cost sharing among multicast receivers. In *Proceedings of 4<sup>th</sup> ACM International Conference on Heterogeneous Networking for Quality, Reliability, Security and Robustness (QShine)*, Vancouver, BC, Canada, August 2007.
- [15] P. Österberg and T. Zhang. Fair cost sharing among multicast receivers. In *Proceedings of 2<sup>nd</sup> IARIA International Conference on Digital Telecommunications (ICDT)*, Silicon Valley, CA, USA, July 2007.