On the Use of QoS Requirements as Base of Multicast Charging[1]

A. Bueno, R. Fabregat, P. Vila, J. Marzo
Institut d'Informàtica i Aplicacions (IIiA), Universitat de Girona (Spain)
e-mail: {bueno, ramon, perev, marzo}@eia.udg.es

Abstract – **In this paper, we propose a charging scheme for cost distribution along a multicast tree when cost is responsibility of the receivers. This scheme focuses on QoS considerations and it does not depend on any specific type of service. The scheme has been designed to be used as a bridge between unicast and multicast services, solving the problem of charging multicast services by means of unicast charging and existing QoS routing mechanisms. We also include a numerical comparison and discussions of the case of non-numerical or relative QoS and on the application to some service examples in order to give a better understanding of the proposal.**

I. INTRODUCTION

Charging for network services is a wide and active subject of study. The main motivations are understanding and/or influencing behaviour, measuring policy compliance, and rational cost allocation/recovery. Charging a point-to-point service is mainly a question of studying the traffic transported, measure it in one way or the other and integrate all this information in a effective but understandable scheme to determine the tariff to be applied to the customer. This is not a simple problem but widely studied.

Charging in ATM networks has been studied by European projects CANCAN and CA$hMAN, mainly for point-to- point services. A further study for ATM and IP networks has been carried out by the European project SUSIE [1].

When point-to-multipoint services are the subject of charging (instead of point-to-point services) additional problems appear. It is not only the amount of connections involved but also the nature of these services: some point-to-multipoint services need additional synchronisation, others may have different categories of users and, in general, most of these types of services have also considerable bandwidth requirements.

In point-to-point services, it is advisable to include in the service description the charge responsible policy [1], (i.e. who will pay for the service: sender, receiver or both and, in this last case, how). In point-to-multipoint it becomes necessary to know how the charge will be distributed because the many variations that can be found: multiple users are involved in these services and, sometimes, they are senders and receivers at the same time. Examples of such services are videoconferencing and commercial information retrieving.

Of the three main scenarios for multicast (one-to-many, many-to-one and many-to-many), this paper focuses in one-to-many services, specifically when charge is assigned to the receivers. In the case of charge being assigned to the sender, there is no cost distribution issue.

Section 2 reviews the concept of *multicast* and the reutilization of data sent through a link. Section 3 discusses the proposal. Section 4 is dedicated to the QoS partitioning problem. Other sections include a numerical example for a video broadcast service, a discussion of the scheme when applied to services where Quality of Service (QoS) is non-numerical or relative, and three examples of the scheme applied to services with different QoS requirements. Conclusions and future work are in Section 8.

II. MULTICAST AND REUTILIZATION OF DATA

The aim of this scheme is to determine how to reflect the resource savings that multicast offers in the cost assigned to every node of a multicast tree.

In this paper we assume that multicast charging is carried out without making any reference to any specific type of service. The cost of each link of the multicast tree is supposed to be determined by the suitable unicast charging scheme for the contracted traffic and QoS.

The idea is to benefit that in point-to-multipoint connections there is the same information travelling by different links. Figure 1a shows the tree that connects a source node A to five destination nodes, from B to F. The same information goes by five different links.

Otherwise, figure 1b shows that the link between A and B (L_{AB}) is used to carry the data to nodes $\{B, \ldots, E\}$. This is possible because of the multicast capacity of node B. If it is cheaper to use L_{BD} than L_{AD} then the resource utilisation is better. In consequence, the cost saved should be distributed between the implied users following some criteria.

The distribution of cost between implied users is not a new problem. Herzog [2] mentions some simple approaches for the case of Internet multicast: Equal Tree Split (ETS), Equal Link Split among Downstream members (ELSD) and Equal Next-Hop Split (ENHS).

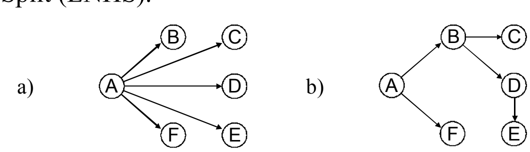

Fig. 1. a) No data reutilization and b) Reutilization of data transported by L_{AB} and L_{BD}

The ETS scheme is the simplest approach to allocating costs. It consists to merely divide the total tree cost equally

[1] This work was partially supported by Universitat de Girona (UdG-DinGruRec2003-GRCT40)

 0-7803-7974-8/03/$17.00

among all receivers. This scheme doesn't discriminate between receivers far from or close to the source and thus doesn't hold receivers accountable for the costs of their individual membership.

The idea that the cost of a particular link is incurred because there is at least one downstream receiver, leads to the ELSD scheme that splits the cost of each link equally among only the downstream receivers.

Finally, the ENHS scheme assigns the cost of a link equally to all the next hop links that are part of the distribution tree. Its motivation is that costs pass on to each downstream next-hop and then the costs of these downstream links are allocated recursively.

The scheme discussed in this paper share some aspects of the last two schemes. A cost distribution along multicast trees is also proposed in [2], but without considering QoS guarantees.

Chuang and Sirbu [3] made an interesting attempt to quantify multicast reduction in the overall network load. The authors empirically found that the cost of a multicast tree varies at a power of the multicast group size. The (normalized) multicast tree cost as is expressed as

$$L_m/L_u = N^k \tag{1}$$

where L_m is the total length of the multicast distribution tree, L_u the average length of unicast routing path, N the multicast group size and *k an* economies of scale factor, ranging between 0 and 1. An extensive validation, with both real and generated networks, shows that the estimation k=0.8 is robust across topological styles and network sizes.

Hurley in [4] proposes a multicast charging scheme based on link weight which is related to DiffServ. Henderson in [5] shows a protocol-independent proposal for multicast pricing, again without QoS guarantees. This paper focuses in QoS considerations.

III. Cost distribution along the multicast tree

Throughout this paper the following notation will be used: The network is represented by a directed graph $G(V, E)$ with cost functions C_L associated to each link $L \in E$. A link from A to B in G is denoted by L_{AB} (or simply AB in subscripts). A multicast session M is denoted as $M=(A,\Delta,Q_S)$ where A is the source node, Δ is a finite set of destination nodes and Q_S are the end-to-end QoS requirements for the connection from A to $D \in \Delta$. For simplicity, we will assume these requirements are equal for all the receivers for a given service S, hence the notation. $T \subseteq G(V, E)$ is a directed tree in G (the multicast tree) that contains all the routes from A to each $D \in \Delta$. Two subsets of Δ are also considered for each $D \in \Delta$: Firstly, P_D is the set of nodes that precede D in its path and secondly, B_D is the set of nodes that share the same parent node with D. Along this paper, it is assumed that T is known and that the cost functions C_L depend on the QoS requirements demanded from L. Only nodes in T will be considered in this discussion.

When the source node A and the first receiver B are connected, there are two possibilities for the addition of a third node C (see fig. 1): connect it either to A or to B. In the first case, node C connects to the tree by A. It is just a couple of 1-1 connections (L_{AB} and L_{AC}) with no multicast benefit. If C connects from B, L_{AB} is used to carry data from A to two nodes, B and C.

The objective now is to decide how to distribute the QoS requirements, and therefore the charge, between the implied nodes. The possibilities range from charging by the complete $L_{AB} + L_{BC}$ connection to C (charging twice by the traffic transported by L_{AB}) to charging to C only by the L_{BC} connection.

Our charging scheme, first proposed in [10], distributes these costs according to the following formulation:

$$\begin{aligned} &Cost(B) = C_{AB}(Q_S) \\ &Cost(C) = \xi\, Cost(B) + \delta_{AB}(Q_S, Q_{AB}) + C_{BC}(Q_{BC}),\ \ \xi \in [0,1] \end{aligned} \tag{2}$$

where Q_{AB} and Q_{BC} are a partition of the QoS requirements Q_S. Q_{AB} and Q_{BC} combined carry the desired service to C with an end-to-end quality of Q_S. These QoS levels are at least as demanding as Q_S, and probably more so, due to the decreasing quality of service of the forwarded traffic.

In our model, node C is responsible for the additional quality required in L_{AB} to achieve a quality Q_S in $L_{AB} + L_{BC}$. The cost increase is noted in the previous formulation as δ_{AB}. The ξ parameter allows balancing between two extreme cases.

When ξ=0 (*fair choice*) the resulting formula is

$$Cost(C) = \delta_{AB}(Q_S, Q_{AB}) + C_{BC}(Q_{BC}) \tag{3}$$

which corresponds to the case when node C is charged strictly by the resources needed to establish its own connection. The main problem with this option arises if intermediate nodes like B can disconnect from the service at any time (dynamic scenario). It would be a problem for the network operator to choose the ξ=0 option and then find that an intermediate node like B in fig. 1b leaves the multicast tree. This situation leaves C, D and E to be charged by the use of L_{AB}. This is against the *predictability* criteria [1] but it's unavoidable if we want to fit the charge to the used resources.

When ξ=1 (*safe choice*), the formula corresponds to a safe option for the network operator, where there is no risk of underestimating the costs of the used resources. But this option is unfair to node C, as the charge does not reflect the real use of resources.

On the positive side, the additional cost δ_{AB}, to ensure the forwarding capability from B, needs to be accounted for just once. When the data stream has enough quality to be forwarded, it can be forwarded as many times as we need it. In fig 1, it would make B as valid a sender as A.

If a new node E is added to the path A-D (as in fig. 1b) its cost is expressed as

$$\begin{aligned} Cost(E) = {} & \xi_B\, Cost(B) + \xi_D\, Cost(D) + \delta_{AB}(Q_S, Q_{AB}'') \\ & + \delta_{BD}(Q_S, Q_{BD}) + C_{DE}(Q_{DE}),\ \ \xi_B, \xi_D \in [0,1] \end{aligned} \tag{4}$$

where Q_{AB}'', Q_{BD} and Q_{DE} are a suitable QoS partition to ensure the service level Q_S from A to E. Note that B is only responsible for the Q_S level through L_{AB}, so $Cost(B)$ remains constant.

In the general case, if we consider not only a unicast connection inside the multicast tree, but a multicast connection as well, the formulation for the scheme is the following:

$$Cost(D) = \sum_{p \in P_D} \xi_{pD} \cdot Cost(p) + \sum_{\substack{p \in P_D \\ b \in B_D}} \lambda_{bD} \partial_{pD}(Q_S, Q_p) + C_{FD}(Q_{FD}) \quad (5)$$

where $D \in \Delta$, P_D is the set of nodes that precedes D in T, F is the last node in P_D (the node D is connected to) and B_D is the set of nodes connected to F. The λ_{bD} allows the increase in cost to be shared among all the nodes in B_D. For this reason, these parameters must fulfill that

$$\sum_{b \in B_D} \lambda_{bD} = 1, \quad \forall D \in \Delta \quad (6)$$

If we evaluate this scheme with the criteria set proposed in [1], a good point is the *usage sensitivity*, as the scheme tries to adjust the charge to the used resources. Another good point is *generality*, as this scheme relay in 1-1 charging connections without compromising with any particular 1-1 charging scheme. This scheme also accomplishes the *co-operative sharing* criteria, due to its multicast nature and it does not seem *easy to fool* because of its simplicity.

The weak point of this scheme is its *predictability*. If the possibility of in-between users disconnecting is taken into account there is always a risk of cost misestimation. If the customer is responsible of the incurred cost when the node where he/she is connected drops, the charge becomes unpredictable from this customer point of view. Otherwise, if the network operator assumes the costs of a disconnecting user and distributes it between all users, a miscalculation in the expected behaviour of users can bring on unforeseen cost or loss of competitiveness. Stability and predictability of prices for end-users has been clearly shown by the INDEX project [7]. Another problem is finding the adequate QoS partition of the end-to-end QoS level ($Q_2 * Q_1^{'} * Q_0^{''} \equiv Q_1 * Q_0^{'} \equiv Q_0$ in fig. 2).

IV. The QoS Partitioning Problem

The need for taking into account the QoS along the whole route is unavoidable if we want to base charging on QoS requirements.

Nevertheless this is not a charging problem. The incidence in cost is not necessarily of great importance. Once the required quality is reached in B to forward the data stream one node further, it can be done as many times as desired. The complexity of the QoS partition problem could be reduced by imposing some limit to the number of hops allowed (some kind of "dispersing limit"), but such a limit goes against the concept of multicast. The service could also be offered at lower quality levels because of distance.

If we put together the complexity of the multicast tree, the diversity of cost functions that can be considered, and the variety of QoS requirements of every receiver, the partition of QoS requirements can be an extremely complex problem. But it is not a new one, as it has been part of the research on the topic of QoS routing for unicast and multicast services. General partitioning problems are simple when dealing with bottleneck requirements, such as bandwidth. However they become intractable for additive (or multiplicative) requirements, such as delay, jitter and loss rate [8].

The problem was first studied [9] in the case of unicast paths. This work showed that there was little difference in performance between simple partitions and optimal ones with certain QoS metrics. But in other cases, e.g. applications that tolerate large packet loss, the difference was significant.

Later, in [10], the problem was studied for call admission of multicast sessions, and a two-phase algorithm was presented. The first phase determines whether there are enough resources to admit the multicast session. After the resource reservation, the second phase releases some of the allocated resources when the characteristics of the receivers and topology of the multicast tree allow it. Two end-to-end QoS division policies were proposed. *Even division policy*, wherever possible, divides equally the end-to-end QoS among the links on a path. *Proportional division policy* balances the load on the links allocating more resources in the less loaded links.

The Constrained Minimum Cost Path and the Constrained Minimum Cost Partition problems (both NP-hard) are considered in [11]. Polynomial ε-approximation algorithms for both problems without assuming any condition on the costs functions are presented.

The first approximated solutions for the case of discrete cost functions were presented in [12], after showing that even the simplest discrete case, i.e. two level cost functions, was NP complete. This case is of particular importance if DiffServ is to be used as a framework for QoS provisioning.

Lorenz and Orda [8] have extensively studied the problem of QoS partitioning as part of QoS routing mechanisms in several papers. They proposed first a fully polynomial ε-approximation to solve the NP-hard Optimally Partitioned Most Probable Path problem for unicast connections. Later they derived ε-approximations for problems OPQR (Optimal Partition of QoS for Routing) and M-OPQ (Multicast Optimal Partition of QoS) for general cost, convex cost (where exact solutions can be achieved) and discrete cost functions and achieved better performance for discrete cost functions than [11] and [12]. In [8] initial results for the complete MOPQ-R problem (QoS partition for multicast routing) are presented.

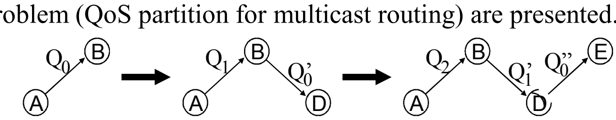

Fig. 2. The QoS partitioning problem

V. Numerical Example

In order to offer a simple context, a simple TV signal (single channel) service is considered in the following numerical example. The QoS could be expressed by two parameters: packet loss and jitter. These two parameters allow expressing the main problems for the cases of video on demand (VoD) or TV signal. At the same time they are simple enough for the intention of this example. Packet loss has a additive quality when expressed in terms of units (packets) and a multiplicative quality when expressed in terms of probability (the case chosen for the example).

The function chosen as the per-second cost function is

$$\mathrm{Cost}(C_i,l,j) = 10^{-5} \cdot \left(\frac{10^{-2}}{l} + \frac{1}{j} \right) \quad (7)$$

where l is the probability of packet loss and j the jitter expressed in seconds. The coefficients are used for giving to l more weight than to j and for scaling the charge to an acceptable amount. The coefficients have been hand-picked to roughly fit Walker's charging proposal [1] example prices for VoD.

This function has the desirable property that variations in different regions have different impact in costs and the drop in cost is not very sudden. Its asymptotical behaviour near the axis corresponds to the intuitive idea of assigning an infinite cost to perfect quality (no packet loss and jitter zero).

If the service quality (Q_S) is defined by $l_S=10^{-4}$ and $j_S=10^{-1}$, a suitable QoS partition for L_{AB}, L_{BC}, L_{CD} and L_{DE} is $l_B= l_D= l_E=10^{-4}/3$, $l_C=10^{-4}/2$ and $j_B=j_D=j_E=10^{-1}/3$, $j_C=10^{-1}/3$. The values obtained of applying (5) are shown in fig. 3 and table 1 for each node. Figure 3 provides a general representation whereas table 1 shows the exact cost values obtained.

Also in fig. 3 and table 1 are the corresponding results for ETS, ELSD and ENHS cost distribution schemes [2]. ELSD and ENHS are equivalent for the simple tree used in this paper. Our scheme fits the charge to used resources, not like ETS, where nodes *B*, *C* and *D* subsidize node *E* traffic. Our scheme also avoids the unfair disparity of charge between nodes *C* and *D* under ELSD and ENHS, where *C* is penalized for not having a node downstream.

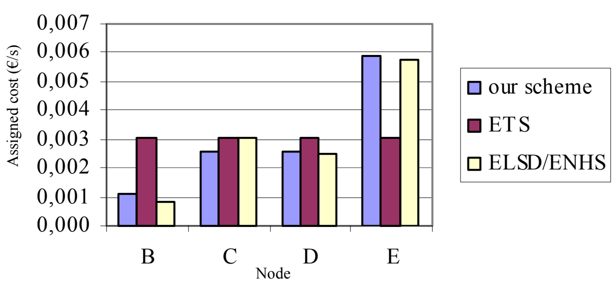

Fig. 3. Cost distribution comparison

TABLE 1
COSTS DISTRIBUTION COMPARISON (€/S)

	OUR SCHEME	ETS	ELSD/ENHS
B	0,0011	0,0030	0,0008
C	0,0026	0,0030	0,0030
D	0,0026	0,0030	0,0025
E	0,0059	0,0030	0,0058

VI. THE QUALITATIVE CASE

The previous example, being numerical, allows a virtually unlimited level of adjustment of quality requirements. Such freedom, while desirable, can be very difficult to manage. Moreover, it can also be not possible in some network architectures. The usual approach to this problem is to use of a limited set of *classes* defined in terms of the QoS offered.

The numerical scenario can be adapted by partitioning the numerical range of every QoS parameter into a set of intervals. Each combination of intervals, one from every QoS parameter, defines a different class. Services are then assigned to the class their specific (numerical) QoS requirements match.

Another way to define service classes is in terms of relative QoS, as in current DiffServ proposals. When specific QoS can not be assured, it may be possible to guarantee a certain degree of quality relative to the network status or between different classes.

The proposed scheme can also be used in this qualitative case. The only requirement is to define a set of classes with an internal order. In this case the options available can be as follow.

When traffic is forwarded in a multicast tree, there is always quality degradation. This degradation makes necessary to increase the QoS in previous path segments (see fig. 2). If the degradation is not significant for the service class, i.e. if although quality is worse, it still complies with the QoS specification of the class, the service can be offered without any quality increase ($\delta_{AB} = 0$) in previous path segments. The scheme formulation becomes simpler in this case:

$$C_C = \xi \,\mathrm{Cost}\,(L_{AB}, Q_S) + \mathrm{Cost}\,(L_{BC}, Q_2), \quad \xi \in [0,1] \quad (8)$$

A second possibility is the originally proposed in (3). If to achieve a given quality level in secondary segments a higher service class is needed (thereof the internal order required in the class set), the increase in cost (δ_{AB}) is the responsibility of the secondary receivers. If the receivers can not assume this cost or if there is not such a higher service class, the service can not be offered with the specified QoS requirements.

A final possibility, depending on the service, is to offer the service at a lower QoS level:

$$C_B = \mathrm{Cost}\,(L_{AB}, Q_S)$$
$$C_C = \xi\, C_B + \mathrm{Cost}\,(L_{BC}, Q_S'), \quad \xi \in [0,1], \quad Q_S' < Q_S \quad (9)$$

This would be only possible if the degraded traffic fits any lower service class offered.

VII. SOME SERVICE EXAMPLES

The first example is a data distribution service. This service would be suitable for a server mirror system. The QoS requirements of this service can be simply to assure no packet loss and a minimum bandwidth.

In this case, the cost of each individual segment is the flat or per time-unit cost of each bandwidth reserve and the cost of assuring no packet loss. The bandwidth requirement does not degrade when traffic is forwarded and the cost of each segment is independent. The only variable parameter is the number of forwarded packets. As one packet lost for one node implies it's also lost for any node connected to it, the additional cost of forwarding lost packets (δ_{AB} in the scheme) is to be charged downstream. δ_{AB} would depend on time (t') and/or the volume of forwarded data (d).

 0-7803-7974-8/03/$17.00

$$C_B = \mathrm{Cost}\ (L_{AB})$$

$$C_C = \xi\ C_B + \delta_{AB}\ (t, d) + \mathrm{Cost}\ (L_{BC}),\quad \xi \in [0,1] \qquad (10)$$

If the network losses too many packets, the cost increase may become prohibitive after a number of hops in the multicast tree.

The second service example is similar to the one used in the numerical example. If we consider a multicast distribution of an audio/video stream, QoS can be expressed in terms of packet loss, jitter and minimum bandwidth. The bandwidth requirement should not degrade when forwarded. Thus the QoS parameters that would require higher quality settings are packet loss (l) and, in a lesser way, jitter (j). This is because the packet loss is much more susceptible of degrading when forwarded.

$$C_B = Cost\ (L_{AB}, l_S, j_S)$$

$$C_C = \xi\ C_B + \delta_{AB}\ (l_S, j_S, l_B, j_B) + Cost\ (L_{BC}, l_C, j_C),\ \xi \in [0,1] \qquad (11)$$

It would be also possible to offer the service with a lower quality level (as in (9)) if packet loss is greater but without rendering the audio/video stream unusable, extending the maximum length of the multicast tree.

The last service example considered is a videoconference service. A possible use of the proposed scheme for such a service could be to consider this service as an aggregate of the previously commented audio/video stream, adding a maximum delay parameter to the QoS specification. This consideration allows integrating sender/receiver, sender-only and receiver-only nodes easily.

The cost distribution scheme operates in this service as in the simple audio/video stream, but in a larger scale. Packet loss and delay limits are the parameters that make necessary higher QoS settings in previous path segments and induce a maximum in the multicast tree branch length.

VIII. CONCLUSIONS AND FUTURE WORK

Along this paper it has been discussed how multicast charging can be based in QoS requirements and how charge can be distributed when the receivers of the multicast connection are responsible of the charge. A fair criterion is to adjust charge to the really used resources. The problem is that to establish an optimum connection tree, the QoS requirements of every segment in the tree must be taken into account.

The strong points of this scheme are its simplicity, its formulation just in terms of QoS requirements and the fact that it relays in unicast charging schemes for each particular segment charging. The problem of QoS partitioning has been thoroughly discussed as well as several QoS interpretations. When service can not be offered at the desired quality has also been discussed.

In reference to future work, the practicability and the implementation of this scheme is our current main concern, as well as an extension for inter-network scenarios. Its interaction with current multicast architectures is a field to explore.

REFERENCES

[1] "Premium IP Services", ACTS Project AC320 (SUSIE) Deliverable 4, document AC320/SUSIE/WP2/*N*/R/L/004/B1, 1999, PDF version found at www.teltec.dcu.ie/susie/

[2] S. Herzog, S. Shenker, and D. Estrin, "Sharing the "cost" of multicast trees: An axiomatic analysis", IEEE/ACM Transactions on Networking, vol. 5, pp. 847-860, Dec. 1997.

[3] J. Chuang, M. Sirbu "Pricing multicast communication: A cost based approach", in Proceedings of INET'98, Geneva, SWITZERLAND, July 1998.

[4] P. Hurley, B. Stiller, T. Braun, "Charging Multicast Communications Based on a Tree Metric", H. Einsiedler, Proceedings of the GI Multicast Workshop'99, Braunschweig, Germany, 20-21 May, 1999

[5] T. Henderson, S. N. Bhatti, "Protocol independent multicast pricing", Proc. NOSSDAV2000 - 10th International Workshop on Network and Operating Systems Support for Digital Audio and Video, The University of North Carolina, USA, 26-28 June 2000

[6] A. Bueno, R. Fabregat, P. Vila "A Quality-Based Cost Distribution Charging Scheme for QoS Multicast Networks", *2nd European Conference on Universal Multiservice Networks, ECUMN'2002*, Colmar, France, april 2002

[7] K. Chu, J. Altmann. "Demand for different qualities of service for internet access: A review of INDEX findings", in *Proc. Royal Society Discussion Meeting on Network mod-elling in the 21st centur*y, London, UK, Dec. 1999.

[8] D. H. Lorenz and A. Orda. "Optimal partition of QoS requirements on unicast paths and multicast trees", INFOCOM'2002, New York , USA, June 2002

[9] R. Nagarajan, J. Kurose, and D. Towsley. Local allocation of end-to-end quality-of-service in high-speed networks. In IFIP TC6 Task Group/WG6.4 International Workshop on Performance of Communication Systems, pages 99-118, Martinique, January 1993.

[10] V. Firoiu and D. Towsley, "Call Admission and Resource Reservation for Multicast Sessions," in IEEE INFOCOM'96, 1996. http://citeseer.nj.nec.com/firoiu96call.html

[11] F. Ergun, R. Sinha, L. Zhang, "QoS Routing with Performance Dependent Costs" Proceedings of the IEEE INFOCOM 2000, Tel-Aviv, Israel, March 2000.

[12] D. Raz and Y. Shavitt. "Optimal partition of QoS requirements with discrete cost functions", in Proceedings of the IEEE INFOCOM 2000, Tel-Aviv, Israel, March 2000.