

A QUALITY-BASED COST DISTRIBUTION CHARGING SCHEME FOR QoS MULTICAST NETWORKS

Antonio Bueno, Ramon Fabregat, Pere Vilà
Institut d'Informàtica i Aplicacions (IiA), Universitat de Girona
Avda. Lluís Santaló, s/n, 17071 Girona (SPAIN)
E-Mail: {bueno, ramon, perev}@eia.udg.es
Phone: +34 972 418 475
Fax: +34 972 418 098

Abstract

This paper presents a new charging scheme for cost distribution along a point-to-multipoint connection when destination nodes are responsible for the cost. The scheme focus on QoS considerations and a complete range of choices is presented. These choices go from a safe scheme for the network operator to a fair scheme to the customer. The in-between cases are also covered. Specific and general problems, like the incidence of users disconnecting dynamically is also discussed.

The aim of this scheme is to encourage the users to disperse the resource demand instead of having a large number of direct connections to the source of the data, which would result in a higher than necessary bandwidth use from the source. This would benefit the overall performance of the network. The implementation of this task must balance between the necessity to offer a competitive service and the risk of not recovering such service cost for the network operator.

Along this paper multicast charging will be made without making any reference to any specific category of service. The proposed scheme is also evaluated with the criteria set proposed in the European ATM charging project CANSAN ([1] and [2]).

1. 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 [3]. 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. Not a simple problem but widely studied.

In ATM networks charging has been studied by European projects CANSAN and CASHMAN, mainly for point-to-point services (see [1] and [2]). Further study for ATM and IP networks has been carried out by European project SUSIE [4].

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 this type of services have also considerable bandwidth requirements.

In point-to-point services, it is advisable to include in service description the charge responsible policy [2], 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.

The three main scenarios for multicast are one-to-many (1-N), many-to-one (N-1) and many-to-many (N-N), as shown in fig. 1. Typical examples are *information distribution* for 1-N, *data collection* for N-1 and videoconference for N-N:




fig. 1. Main scenarios for multicast

The main subject of discussion in this paper is the 1-N service. The N-1 service is not profitable for this paper due to the unicity of the information sent by each channel. Finally, in this paper, the N-N service will be considered as an aggregate of point-to-multipoint connections.

This paper is therefore focused in 1-N services, specifically in charging based in QoS parameters when charge is assigned to the receivers. If charge is assigned to the sender, there is no cost distribution




fig. 2. Working scenario

issue. The network operator should decide and charge the sender by the sum of all connections made.

Section 2 of this paper reviews the concept of *multicast* and the reutilization of data sent through a link. Section 3 presents the proposal: first, two extreme cases: a no-risks choice and a fair, more competitive but risky choice; then an integration of both covering all the in-between possibilities and their advantages and drawbacks. Section 4 offers a numerical example. Finally, conclusions and future work are presented.

2. Multicast and reutilization of data

One of the main characteristics of multicast is that duplication of packets (*cells*, in ATM terminology) only happens when routing paths diverge. By avoiding the transmission of duplicated cells over any link, significant bandwidth savings can be achieved. This is why a point-to-multipoint connection is a tree of links, with the source as the root, instead of a set of point-to-point connections.

The aim of this proposal is to make use of the diversity of possible configurations of the multicast tree to implement a discount system to encourage the users to disperse the resource demand, profiting by the capability of charging to shape the resource use. To distribute the connections along the tree rather than concentrate the connections near the source would benefit the overall performance of the network. If this objective were achieved, route selection considerations would be included in the charge. Therefore charge can be considered as the route selection parameter. This is a very straightforward and desirable option because different routes are presented to the customer who choose just in terms of charge.

Along this paper multicast charging is carried out without making any reference to any specific category of service. The cost of each link of the multicast tree is supposed to be determined with the suitable unicast charging scheme for the contracted traffic and quality of service (QoS).

As mentioned before, the idea is to benefit from the fact that in point-to-multipoint connections there is the same information travelling by different links. Figure 3a shows the tree established to connect a source node A to five destination nodes, B, C, D, E and F. The same information goes by five different links.




fig. 3a. No data reutilization




fig 3b. Reutilization of data sent by L_{AB} and L_{BD}

Otherwise, figure 3b shows that the link between A and B (L_{AB}) is used to carry the data to nodes B to 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. But how should this distribution be made? How distribute these savings?

The distribution of cost between implied users is not a new problem, nor a specific ATM problem. [5] 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). In [6] a cost distribution along multicast trees is proposed, but without QoS guarantees.

The ETS scheme is the simplest approach to allocating costs. It consists to merely divide the total tree cost equally among all receivers. This scheme does not discriminate between receivers far from or close to the source and thus does not 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 the idea that costs pass on to each downstream next-hop and then the costs of those downstream links are allocated recursively.

The scheme presented in this paper share some aspects of the last two schemes, focusing in QoS considerations.

An interesting attempt to quantify multicast's reduction in overall network load was made by Chuang and Sirbu [7]. 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 is expressed in [7] as

$$L_m/L_u = N^k$$


where L_m is the total length of 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. Extensive validation with both real and generated networks show that the estimation $k=0.8$ is robust across topological styles and network sizes.

In [8] a multicast charging scheme based on link weight is proposed and related to DiffServ. [9] shows a protocol-independent proposal for multicast pricing, again without QoS guarantees.

3. Cost distribution along the multicast tree

When building the multicast tree, ITU-T and ATM Forum User-Network Interface specifications ([10] and [11]) require setting up the destination nodes sequentially. First, a point-to-point connection is set-up between the source node and the first destination node. After this set-up is

complete, additional destination nodes are added sequentially in the same manner. The scenario of this discussion is when the first connection between source node A and first receiver B has been established. Then there are two possibilities to connect a third node C: to connect to node A or to node B.



In the first case, node C connects to the tree from A (fig. 4a). It is just a couple of 1-1 connections with no multicast benefit. If C connects from B (fig. 4b), L_{AB} is used to carry data to two nodes, B and C.

The idea is to offer the customer a better connection choice, in terms of charge. The customer would get the desired service at a better price and the resource demand is not concentrated around the source of the multicast tree. The objective is then to decide how to distribute the charge between the implied nodes.

The range of possibilities goes from charge by the complete $L_{AB}+L_{BC}$ connection to C (charging twice by the traffic sent by L_{AB}) to charge to C only by the L_{BC} connection. An alternative option is any compromise between these extreme cases. In this case it is also possible to contemplate some discount policy applicable to node B due to the L_{AB} sharing.

Safe choice for the network operator

It would be unfair to charge to C by the whole $L_{AB}+L_{BC}$ connection, as it would be charging twice by L_{AB} (one to B and another to C). The only justification to choose this option would be the simplicity of management, but such an option would only be chosen if direct A-C connection (some L_{AC}) has a cost higher than the full $L_{AB}+L_{BC}$ path. The charge C_B to node B and C_C to node C have the following expression:

$$C_B = \text{Cost}(L_{AB}, Q_0)$$

$$C_C = \text{Cost}(L_{AB}, Q_1) + \text{Cost}(L_{BC}, Q_2)$$

Note that three different levels of QoS, Q_0 , Q_1 and Q_2 appear. This is because the cost of L_{BC} is not the only to be accounted, but also the additional costs induced in L_{AB} to assure the agreed QoS. This would mean that in C_C the required quality over L_{AB} is higher. When a connection through L_{AB} is established, a certain QoS grade Q_0 (associated to the service) is guaranteed. This grade is enough to offer the quality required to B, but not necessarily enough to carry the data stream to C with quality Q_0 . In consequence, to assure data arrives to C with the required quality level (Q_0) additional quality, Q_1 (better than Q_0) needs to be guaranteed over L_{AB} . The cost of L_{AB} can then be expressed with the following formula:

$$\text{Cost}(L_{AB}, Q_1) = \text{Cost}(L_{AB}, Q_0) + \delta_{AB}(Q_0, Q_1)$$

where δ_{AB} is the cost of the additional quality over L_{AB} .

Furthermore, a Q_0 QoS level is not enough for L_{BC} , and a new quality level Q_2 (better than Q_0) is needed. The combination of Q_1 and Q_2 should give then the expected quality level Q_0 . Therefore the previous formulae become:

$$C_B = \text{Cost}(L_{AB}, Q_0)$$

$$C_C = \text{Cost}(L_{AB}, Q_1) + \text{Cost}(L_{BC}, Q_2) = \text{Cost}(L_{AB}, Q_0) + \delta_{AB}(Q_0, Q_1) + \text{Cost}(L_{BC}, Q_2)$$

and the total cost allocated is

$$\sum C_i = C_B + C_C = 2 \cdot \text{Cost}(L_{AB}, Q_0) + \delta_{AB}(Q_0, Q_1) + \text{Cost}(L_{BC}, Q_2)$$

It is clear in the formula that the same data is charged twice. One advantage of this option is that A-B connection and the A-C connection can be independently managed. Another advantage is that network operator does not need to worry about in-between users disconnecting. Users disconnecting can be a serious problem if C_C charge is fixed counting on that node B is being charged an amount C_B . On leaving, node B leaves node C as full responsible for the L_{AB} cost. The drawback of this option is that the A-B-C connection is not being charged by the real resources used and it would only be chosen if direct A-C connection is more expensive or it does not exist.




fig. 5a. L_{AB} accumulated charge vs. time




Fig. 5b L_{BC} charge vs. time

Fair choice for the customer

On the other hand, to charge node C just by its own added costs, the charges C_B and C_C would have then the following expressions:

$$C_B = \text{Cost}(L_{AB}, Q_0)$$

$$C_C = \text{Cost}(L_{BC}, Q_2) + \delta_{AB}(Q_0, Q_1)$$

In this case, the total cost allocated is

$$\sum C_i = C_B + C_C = \text{Cost}(L_{AB}, Q_0) + \delta_{AB}(Q_0, Q_1) + \text{Cost}(L_{BC}, Q_2)$$

and corresponds to the used resources. In figs. 6a and 6b the cost of each link, when nodes B and C connect to the multicast tree, is represented graphically:




fig. 6a. L_{AB} accumulated charge vs. time




fig. 6b. L_{BC} charge vs. time

In figs. 6a and 6b the functions represent the charge by the time interval $[0, t]$ corresponding to each link. In these figures charge is assumed to be proportional to time. This assumption corresponds to a unicast charging scheme where all traffic characteristic are summarised in just one parameter (the slope in the graphic) and then charge is done by time. Examples of such schemes are Kelly's and Envelope charging schemes ([12], [13] and [1]), included in CANCAN documentation.

The necessity of taking into account the QoS along the whole route is one of the main drawbacks of this approach. Just think in the case of node E in fig. 2b. No matter how far is a node from the source. Every quality level, from the source to the destination, has to be taken into account, as shown in fig. 7.

Nevertheless this is a management problem, not a charging one. The incidence in cost is not necessarily of great importance. Once we have the required quality in B to represent the data stream one further node, it can be done as many times as we want. The quality backtracking problem

could be reduced by imposing some limit to the number of hops allowed (some kind of "dispersing limit"), but this would be against the aim of this proposal since it limits the connection choices available for the customer.




fig. 7. Quality backtracking problem

Integration of both cases

We can integrate both extreme cases and cover the in-between cases with the following expression:

$$C_C = \xi \cdot \text{Cost}(L_{AB}, Q_0) + \delta_{AB}(Q_0, Q_1) + \text{Cost}(L_{BC}, Q_2) \text{ where } \xi \in [0, 1]$$

If we assume that $\xi=1$, we get the first formula which corresponds to the secure option for the network operator, as there is no risk of underestimate the costs of the used resources. On the contrary this option is unfair to the node C, as charge does not reflects the real use of resources.

When $\xi=0$, the obtained formula is $C_C = \delta_{AB}(Q_0, Q_1) + \text{Cost}(L_{BC}, Q_0)$, which corresponds to the case when node C is only charged by the additional resources needed to establish its connection. The first problem that arises with this choice is when we consider the possibility that the intermediate node B can disconnect from the service at any time. It would be a problem for the network operator to choose the $\xi=0$ option and then found that B leaves the multicast tree. And to assign to C the full charge depending on the behaviour of B is against any fairness or usage sensitivity criteria [1].




fig. 8. L_{BC} charge varies with ξ

Both extreme options can be pondered by choosing a $\xi \in]0, 1[$ (see fig. 8), but to determine the optimum value would require considerable knowledge about the behaviour of such services. And with such scheme there would be always some risk of undercharge to the network operator.

Note that the only adverse disconnections are those from intermediate nodes (the *branches*, not the *leaves* of the tree). In the situation of fig. 3b, such nodes are B and D. If B disconnects it leaves C, D and E to be charged by the use of L_{AB} . Observe that, possibly against intuition, the more nodes connected to such a disconnecting node, the lower is

the cost impact of the node disconnection because this cost can be shared between them.

On the positive side, we have that the additional cost δ_{AB} to ensure the re-sending capability from B needs to be accounted just once. When the data stream have enough quality to be re-sent, it can be re-sent as many times as we need it. In the example used along all this discussion, it would made B a sender as valid as A.

A final possibility to take into account is to include some discount to the node B to compensate any tendency to favour the secondary nodes. The discounted amount should be shared between the secondary nodes (just C, in the example). The correspondent expressions would be:

$$C_B = \lambda \cdot Cost(L_{AB}, Q_0)$$

$$C_C = (\xi + 1 - \lambda) \cdot Cost(L_{AB}, Q_0)$$

$$+ \delta_{AB}(Q_0, Q_1) + Cost(L_{BC}, Q_2)$$

where $\lambda, \xi \in [0, 1]$

If we evaluate this scheme with the criteria set proposed in CANSAN ([1] and [2], a good point is *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 accomplish the *co-operative sharing* criteria, due to its multicast nature and it does not seem *easy to fool* because of its simplicity.

The drawbacks of this scheme are its *fairness* and *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 made responsible of incurred cost when the node where he 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 are highly desirable, as the INDEX project [14] has clearly shown.

After all this discussion, it seems clear that any attempt to adjust the charge to the additional resources bring about some risk to the customer, the network operator or both of them. If the main objective is to avoid unforeseen cost to the network operator there is a safer option, a less customer-oriented scenario: all decisions about routing are left to the network operator, who reflects saved resources in *a-posteriori* discounts to customers. This would make easier the determination of saved resources, especially if users can disconnect in any moment, and would avoid any loss risk to the network operator. The main drawback would be the difficulty to customers to accept such unpredictable scheme.

4. Numerical example

This section tries to clarify the formal formulation of the proposal with a numerical example. In order to offer a simple context a simple TV signal (single channel) service is considered. The QoS could be expressed by two parameters: packet loss and jitter¹.

These two parameters allow to express well enough the main problems for the case of video on demand or TV signal. At the same time they are simple enough for the intention of this example. *Packet loss* is preferred to *packets arrived* for its simplest formulation. It has also an 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).

For the numerical example, some functions were considered. All of them are decreasing functions expressed in the previously mentioned parameters: A lineal function $f(x)=ax+b$, $a < 0$ and a negative exponential function where discarded in the last round. The chosen function is a function inversely proportional to the QoS parameters.

The lineal function is the simplest option. The problem is that it gives the same importance to a variation of quality without taking into account if it's already a high or low quality.

The negative exponential function doesn't have this problem. It decreases very quickly and this is consistent with QoS parameters costs. The main drawback is the low impact of variations of near-zero QoS parameters (when expressed in the terms of this paper).

Finally, the inversely proportional function solves, at least partially, both problems. Variations in different regions have different impact in costs and the drop in costs is not as sudden as with the neg. exp. function. The main characteristic is its asymptotical behaviour near the axis. Service characteristics and/or physical network constraints naturally avoid this function singularity.

In any case, any of these three functions should be truncated to an interval $[\epsilon, c]$ due to service or network limitations/characteristics. The lower limit would correspond to limits in the network capacity. The upper limit would answer to admissible quality limits for the service. This truncation is not included in the presented numerical example.

In this example, the charge for both parameters will be a function inversely proportional to both parameters. Such a function corresponds to the intuitive idea of assigning an infinite cost to perfect quality (in our example, no packet loss and jitter zero).

¹ It is defined as the absolute difference between two adjacent packets arrival times minus their departure times. A significant percentile is usually taken as a numerical representative. See RFC 2598 for an example.

For this example we use the expression $charge/sec. = 10^{-5} \cdot \left(\frac{10^{-2}}{p} + \frac{1}{s} \right)$ where p is the probability of packet loss and s the jitter expressed in seconds. The coefficients are used to give p more weight than to s and to scale the charge to an acceptable amount. For example, if the parameters for a high quality video stream are $p=10^{-4}$ and $s=10^{-1}$, the charge is around €4 per hour (0.0011 €/s). The parameters have been hand-picked to roughly fit Walker's charging proposal [1] example prices for VoD.

In addition to these coefficients, this formulation assumes the contribution of probability of packet loss and jitter can (or should) be aggregated by addition. The only justifications are simplicity and to avoid the extreme variations of the multiplicative alternative $((1/p) \cdot (1/s))$.

If these are the parameters for our example, i.e. the initial quality (Q_0) is given by $p=10^{-4}$ and $s=10^{-1}$, the qualities Q_1 and Q_2 for L_{AB} can be achieved with $p=0.5 \cdot 10^{-4}$ and $s=0.5 \cdot 10^{-1}$. The values for section 2 formulas are:

$$Cost(L_{AB}, Q_0) = 0.0011 \text{ €/s}$$

$$Cost(L_{AB}, Q_1) = 0.0022 \text{ €/s}$$

$$Cost(L_{BC}, Q_2) = 0.0022 \text{ €/s}$$

| | Safe choice | Fair Choice | Integration of both cases |
|-------|-------------|-------------|--|
| C_B | 0.0011 €/s | 0.0011 €/s | 0.0011 €/s |
| C_C | 0.0044 €/s | 0.0033 €/s | anything between 0.0033 €/s and 0.0044 €/s |

Note that the increase in cost of C_C is to be shared between the tree nodes connected to B .

Similar considerations can be made for the case where quality is not determined by quantitative parameters, but by a set of quality classes. This can be seen as a discretisation of the quantitative case. In the qualitative case the availability of a class that fulfil the service QoS requirements is probably an important point.

5. Conclusions and future work

Along this paper it has been discussed how charge can be distributed when the receivers of a point-to-multipoint connection are responsible of the charge. The aim was to define a charging policy to encourage users to disperse resource demand along the network instead of concentrate the demand in the source node of the multicast tree.

To transport data through more than one link implies that higher QoS levels need to be assured and consequently, charged. A fair criterion is to adjust charge to the really used resources. The problem is that to establish

sequentially an optimum connection tree most (if not all) of the previous tree must be taken into account.

There is a safe scheme to the network operator: to charge every user by the whole path from the source. But this option can be very punitive to users and it goes against the dispersing resources target. Any other fairer possibility implies some degree of loss risk to the network operator or undesirable surprise to the customer.

Another possibility for this scheme is to be used by the network operator to evaluate the tree cost and use it to determine a more marketable tariff. This could be suitable for a non-commercial (i.e. academical) scenario.

The major problem is well known in multicast. If the possibility of in-between users disconnecting dynamically from the multicast tree is included, then any solution other than to charge every user by the whole path is compromised. The charge distribution should be recalculated and an optimal solution could become the worst.

Finally note that all this discussion is independent on the traffic category (CBR, VBR, ... in the ATM case) of the connections. But some of the services offered through this categories (especially those with delay boundaries) may have something to say about this discussion because their QoS requirements could not be satisfied if data is re-sent. In general, the limitations in the QoS parameters induce a maximum in the number of hops that can be made to carry the data.

In reference to future work, the applicability and practicability of this scheme is our current main concern, as well as an extension for inter-network scenarios.

Another subject of further study is N-N services. The assumption of considering any N-N connection as an aggregation of 1-N connection imposes some unnecessary restrictions. N-N services are more flexible and the development of a similar scheme to the proposed in this paper can be studied.

6. Acknowledgments

This work was partially supported by the Spanish Research Council (CICYT) under contract TEL-99-0976. The authors would like to thank the members of the BCDS group at the University of Girona for their useful comments.

7. References

- [1]. "ATM Charging Schemes (Review of ATM Charging Schemes and Performance Issues)", ACTS Project AC014 (CANCAN) Deliverable 5, document AC014/QMW/-/DS./P/005/B1, August 1996. PDF version found at <http://gryphon.elec.qmw.ac.uk/staffinfo/eric/cancan/>

- [2] "Final Report on Static Charging Schemes and their performance", ACTS Project AC014 (CANCAN) Deliverable 9a, document AC014/QMW/DS/P/350/a7, November 1997. PDF version found at <http://gryphon.elec.qmw.ac.uk/staffinfo/eric/cancan/>
- [3] "Internet accounting: Background", C. Mills, D. Hirsh, G. Ruth, RFC1272, 1991
- [4] "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/
- [5] "Accounting and Access Control for Multicast Distributions: Models and Mechanisms", Shai Herzog, PhD dissertation, University of Southern California, August 1996, found at <http://netweb.usc.edu/herzog/docs/thesis.ps>
- [6] "Sharing the "cost" of multicast trees: An axiomatic analysis", Shai Herzog, Scott Shenker, and Deborah Estrin, IEEE/ACM Transactions on Networking, vol. 5, pp. 847-860, Dec. 1997.
- [7] "Pricing multicast communication: A cost based approach", J. Chuang and M. Sirbu, in Proceedings of INET'98, Geneva, SWITZERLAND, July 1998.
- [8] "Charging Multicast Communications Based on a Tree Metric", H. Einsiedler, P. Hurley, B. Stiller, T. Braun, Proceedings of the GI Multicast Workshop'99, Braunschweig, Germany, 20-21 May, 1999
- [9] "Protocol independent multicast pricing", T. Henderson, S. N. Bhatti, 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
- [10] "Broadband Integrated Service Digital Network (B-ISDN) Digital Subscriber Signaling System No. 2 (DSS2) User Network Interface Layer 3 Specification for Basic Call/Connection Control ITU-T Recommendation Q2931", International Telecommunication Union, Geneva, 1994
- [11] "ATM User-Network Interface Specification Version 3.1", ATM Forum, July 21, 1994
- [12] "Tariffs, Policing and Admission Control for Multiservice Networks", Frank P. Kelly, 10th UK Teletraffic Symposium, BT Laboratories, Martlesham Heath, April 1993 and CANCAN, ATM Charging, Background paper 1.7
- [13] "Charging and Accounting for Bursty Connections", Frank P. Kelly, Statistical Laboratory, University of Cambridge, 16 Mill Lane, Cambridge CB2 1SB, England, July 1993 and CANCAN Background paper 1.15
- [14] K. Chu and J. Altmann. Demand for different qualities of service for internet access: A review of INDEX findings. In *Proc. Royal Society Discussion Meeting on Network modelling in the 21st century*, London, UK, Dec. 1999.