Universität Stuttgart

INSTITUT FÜR NACHRICHTENVERMITTLUNG UND DATENVERARBEITUNG
Prof. Dr.-Ing. Dr. h. c. P. J. Kühn

Interner Bericht / Internal Report **Nº 29**

Titel / Title

Charging in Multi-Service Networks

Verfasser / Author(s)

Stefan Bodamer

Datum / Date 12.11.98

Umfang / Size 15 Seiten / Pages

Quelle / Source Report of a small project sponsored by Siemens ICN M RP 11

Schlüsselworte / Keywords Internet, ATM, Charging, Overview, Effective Bandwidth, Priority

Beitrag der Arbeit / Achievement

Review of proposed architectures and mechanisms for charging and pricing in IP and ATM networks

Kurzfassung / Abstract

Charging is becoming a more and more interesting issue in the Internet as well as in ATM networks. This report first gives an overview on pricing schemes including a rough classification. Also some problems associated especially with Internet pricing are addressed. The most important concepts for pricing in both types of networks are presented in a short form. The advantages and disadvantages of each scheme are pointed out. Finally, the pricing approaches are compared in order to be able to decide which are the most promising solutions.

1 Introduction

With the Internet becoming more and more popular not only technical but also economical aspects are a matter of increasing interest. Among the members of the Internet community different expectations to Internet economics can be observed.

Internet users want to get best service for their money. While there are various types of applications with completely different characteristics and requirements, not only some kind of best-effort service but also the possibility to choose between different levels of quality of service is a desirable feature from the users' viewpoint.

Service providers compete to keep a good market position and to make revenue from the network. Therefore they have to offer services attractive in both quality and price on the one hand, while on the other hand limiting their network costs. Appropriate pricing schemes have to be developed to achieve these goals.

As a consequence from this also equipment providers are influenced by economical aspects. New pricing schemes raise the need for new architectures and mechanisms which have to be implemented in routers or edge devices. Also accounting and billing impose questions concerning the software needed to operate a network running a certain pricing method.

While pricing has a long tradition in classical telecommunication networks like the telephone network a remarkable characteristic of the Internet is that service has been for free. Free means that only access fees have to be paid on a per month or per year basis. This kind of flat rate pricing, however, does not provide an appropriate instrument to cope with increasing congestion in the Internet.

Therefore another pricing model becomes more and more interesting: usage-sensitive or congestion pricing. In addition to the flat access price a user pays a variable charge for each byte sent over the network. The basic idea is further that this charge per unit of bandwidth is higher in times and locations of congestion, while it may be zero if there is no congestion.

A third pricing model is called transaction-based pricing in the literature [33, 41]. This type of charging scheme also takes the usage of network resources into account. However, not the amount of data sent over the network determines the price but the characteristics of a transaction (or connection). So transaction pricing mainly refers to connection-oriented transmission like in the telephone system or in ATM networks.

The next section gives an overview to Internet charging and presents some interesting proposals for charging schemes. The introduction of pricing in the Internet is also influenced by research in the ATM area, as both running IP over ATM as well as using new resource reservation methods like RSVP in the Internet are widely discussed options. Therefore some approaches for charging in ATM networks are described in Section 3. In the last section of this report some comparison and validation of charging schemes is done. The references section does not only contain papers referenced in the report but also further interesting sources of information.

2 Internet Charging

2.1 Overview (MacKie-Mason and Varian)

In [28] MacKie-Mason and Varian give answers to important questions associated with Internet economics.

2.1.1 Current Ways of Internet Pricing

The most important pricing mechanisms in today‘s Internet is connection or flat rate pricing. A user pays an annual or monthly fee for being connected to the Internet. Usage of network resources is only limited by the access rate depending on the access technology (e. g. POTS, ISDN, ADSL).

A subcase of connection pricing is denoted as committed information rate pricing in [28]. This mechanism sometimes applied in private networks provides a guaranteed rate for which an additional flat price has to be paid.

Moreover, some providers use special forms of pricing, e.g. a separate charge for e-mails or charging for access to an international link [3].

2.1.2 Prices and Costs

MacKie-Mason and Varian point out that according to economic theory prices should always reflect costs. Network costs can be divided into three parts:

- connection costs
- capacity costs
- social costs due to congestion

While connection costs can be adequately covered by connection pricing using a flat rate, capacity costs and social costs require a usage-based pricing scheme.

2.1.3 Pricing and Congestion Control

Congestion in the Internet is a problem which is more serious than in classical networks due to the wide range of usage rates. Prohibition of resource intensive applications can only be a temporary solution. The network only sees IP packets and even the field in the IP header indicating the transport protocol (e. g. TCP, UDP) may be manipulated by the user or the application.

One can think of various solutions to the congestion problem:

- **Overdimensioning**: This is a sort of general solution which, however, may become very costly. The authors of [28] expect that the user requirements increase with the same speed as the network resources. So there will always be a scarcity of resources.
- **Resource reservation**: Reservation mechanisms like RSVP that guarantee a certain bandwidth maintaining a specified delay. However, there are many objections against an introduction of such protocols in the Internet community.
- **Voluntary control**: There are various possible forms of this kind of congestion control mechanism whose effectiveness, however, might be low especially since the number of inexperienced users increases:

- doing without a costly application (e. g. high bitrate video communication) in times of congestion
- indicating the importance of information setting a specific value in the service type field of the IP header
- slow start mechanism in TCP

Usage based pricing - in contrast to voluntary control - gives an incentive to the users to reduce traffic or to choose the right level of priority in congestion situations. So users have the possibility to express the value which they associate to their traffic. Furthermore, pricing gives an additional indication to the network operator how serious congestion is and therefore the right signal to increase the capacity.

2.1.4 Problems with Internet Pricing

There are also some problems associated with Internet pricing, that are widely unsolved until now:

- sender-receiver problem, i.e. the question who pays for data transfer and how to realize a concept where the receiver pays
- complexity of packet counting
- identification of users
- support by appropriate router architectures

2.2 Responsive Pricing (MacKie-Mason, Murphy and Murphy)

In [27, 29, 30] the authors present a pricing scheme called responsive pricing where feedback signals between network and user play a major role.

The quality of a network is mainly determined by user satisfaction. This satisfaction is only roughly represented by engineering measures like loss ratio or delay. The reason for this is that there are great differences in the requirements depending of the types of users and applications:

- Real-time applications are less loss sensitive than e. g. remote login applications.
- Interactive applications suffer from delay in contrast to file transfer applications
- A certain level of QoS may be valuated differently by different users even if they use the same application.
- Some users require deterministic or at least statistic guarantees while others only need a best effort service.

These examples show that a separation of traffic according to only some few classes of applications is not a solution to the problem. The consequence from this is that the user has to be part of the network control loop. In the process of evaluating the feedback from the network the user can be supported by intelligence put into his workstation.

Feedback signals given by the network can be evaluated by the user on different time scales:

- On a long-term basis users decide whether they use a network service or not.
- Users choose the time of day when they are on-line trying to avoid phases of congestion.

- Within a session users adapt their traffic intensity or their QoS requirements depending on the state of the network.
- On the packet level congestion control mechanisms are provided by protocols like TCP.

The user may, however, ignore feedback signals (e. g. by modifying TCP) and so increase congestion and social network costs (see Section 2.1). Therefore pricing is suggested as a form of feedback which gives an incentive to the user to react to signals coming from the network.

As a simple approach to realize this a closed-loop feedback scheme is proposed in [27]. The state of the network is measured, e. g. in terms of buffer occupancy in the routers. From this state a price per packet is derived which is signalled to the users in regular intervals. A problem is, however, to find an appropriate relation between congestion state and price. Furthermore, a detailed mechanism and the time interval have to be defined for the distribution of prices.

A more enhanced method to use prices as a form of feedback is the smart market approach. When the network is in congestion state the user can indicate the value of a packet by associating a “bid” to the packet. This bid is the price the user is willing to pay for the transmission of this packet. In a congested router, only a certain amount of packets is accepted. This is achieved by sorting the bids giving access only to packets with higher bids and discarding the others. All packets being accepted pay the same price which is equal to the lowest bid of any accepted packet. This mechanism corresponds to the concept of a Vickrey auction.

The appealing property of this scheme is that the users can exactly valuate the transmission of each packet. MacKie-Mason et al., however, identify some problems with the smart market approach:

- The packets arrive continuously. The acceptance decision, however, has to be made on a time-slice basis.
- Users don‘t know in advance how much they have to pay.
- The smart market does not provide any guarantees. Only relative priorities can be indicated.
- The way of implementing a smart market in a real network with many possibly congested routers is unknown. Especially the questions how to solve the sender-receiver problem addressed in 2.1.4 is open.

Both realizations, the closed loop feedback scheme and the smart market approach, are intended to set prices only in the congestion case.

2.3 Priority Pricing (Gupta et al., Cocchi et al.)

Like the smart market approach described in Section 2.2 the pricing scheme presented by Gupta et al. in [18] uses priorities on the packet level. The authors propose to differentiate Internet traffic according to delay and loss requirements. As an example a separation into four levels of priority is given (with lower values indicating higher priority):

- Priority 1: Real-time services that are loss-sensitive
- Priority 2: Real-time services which are not loss-sensitive
- Priority 3: High priority best-effort traffic
- Priority 4: Low priority best-effort traffic

The priority of a packet can be indicated using the type of service (TOS) field in the IP header. Inside the network the authors suggest to use “priority queues” in the routers without specifying, however, to which kind of priority (space or delay priority) they refer. The assignment of prices to the priorities is left open to the ISP.

A similar charging scheme is presented in an early paper by Cocchi et al. [6]. Two bits in the IP header are used to indicate whether a packet is loss-sensitive or delay-sensitive, respectively. The price per packet is $k + 1$ times a base price with k denoting the number of priority bits set.

2.4 Expected Capacity Allocation (Clark)

In [5] pricing is introduced as a means to offer different options in network usage to the users in a controlled manner. In the current Internet there is no possibility to obtain better quality by paying more which can become annoying in congestion situations. In this context Clark points out that congestion can mainly be measured in form of reduced throughput leading to transmission delay on the application level (e. g. single character in remote login, text page or picture in WWW, file in FTP), while packet level delay is not a major problem. This observation, however, is restricted to data services based on TCP.

Concerning the realization of pricing in the Internet Clark proposes a charging scheme called expected capacity allocation. This approach can be understood as a compromise between guaranteed and best-effort service. User and network provider negotiate a sort of traffic contract containing an expected capacity profile. At the network ingress packets are tagged as being “in” or “out” with respect to the capacity profile. This task is done by a token bucket. The token bucket may be realized in the same way as a leaky bucket for ATM/VBR traffic monitoring a cell stream with regard to sustainable rate and maximum burst size. Inside the network “in” and “out” packets are handled differently in such a way that “out” packets are discarded first in congested routers. The detailed form of the capacity profile is determined by the ISP as well as the manner how “in” and “out” packets are measured and treated in the network.

The expected capacity allocation scheme is strongly concatenated with the upcoming diffrentiated services (DiffServ) approach of providing QoS in the Internet. On the other hand it also reminds of traffic management in ATM with UPC monitoring traffic according to a contract and using the CLP bit in the cell header. However, there are mainly two differences between the approaches:

- The traffic contract does not refer to a connection like in ATM but holds on a long-term basis. In this sense expected capacity pricing is more like flat rate pricing where the user pays for a certain access rate. The user has, e. g., to pay for the rate and the size of the token bucket. The bucket values, however, are no upper usage bounds but only thresholds for differentiating packets in the congestion case.
- The expected capacity allocation method does not give any guarantees with respect to throughput or delay. The only guarantee that is given is that the probability for a packet being discarded along a certain path within the network is lower for “in” packets than for “out” packets. A consequence from this is that unlike in ATM no CAC is required to maintain any guarantees.

The author also addresses the sender-receiver problem mentioned in Section 2.1.4. Especially for WWW traffic a “receiver pays” model is useful as the sender, which is a server providing public information goods in many cases, has no major incentive to increase throughput by paying larger fees. As a simple solution to that problem Clark suggests to additionally perform a conformance test at the sender with respect to the receiver profile. This approach, however, has some shortcomings:

- The receiver profile is unknown at the sender location unless a complex signalling mechanism for the exchange of profiles is installed.
- Sender and receiver may belong to different providers using different types of capacity profiles.
- In a multicast session the sender may even not know all receivers.

To overcome these drawbacks a more enhanced approach to handle the problem is presented. In this approach packets for which the receiver has to pay are tagged at congested routers and possibly discarded at the receiver site according to the receiver profile. This means that all packets are transmitted over the network but discarding them at the receiver bucket gives the right signal for TCP to reduce the rate. Questions like how to decide whether the sender or the receiver profile have to be applied, however, still remain open.

2.5 Paris Metro Pricing (Odlyzko)

A pricing scheme called PMP (Paris Metro Pricing) has been developed at AT&T Labs [37, 38, 39]. The name of the approach indicates that it was inspired by the Paris Metro system. Until some years ago the Paris Metro offered transportation in first and second class cars with identical number and quality of seats but with different prices for the tickets. The system operated in an equilibrium with the first class being less crowded but more expensive.

The analogy between PMP and the Paris Metro is that there are several distinct networks with different prices in the PMP approach. Distinct in this context means that on each link a fixed share of the capacity is associated to each traffic class. The quality of service will be higher in the more expensive subnetworks as the number of users is smaller. This is the same effect as that being observed when regarding the performance in networks of different providers with different access fees. The advantage of this approach is simplicity due to its self-regulating nature.

Instead of using separate networks it would also be possible to use static priorities (as proposed in [18]) or different weights if a fair queueing scheduling technique (e. g. WFQ) is applied. One could also think of a mixture of different methods.

An implementation of PMP could use the priority field available in the IP header restricting the number of priority levels to 8 (3 bits). Odlyzko suggests to limit the number of priorities to 4. The IP priority bit could be set on the workstation by a “wrapper” software handling all IP traffic on that machine. Inside the network some changes are required in the routers to install some kind of priority service. However, this could be done step by step thereby easing migration to a new network structure. Furthermore, accounting hardware and software is only required in the edge routers but not in the network core.

Issues like the “receiver pays” problem or support of multicast services is left open for further research in the papers.

3 Charging in ATM Networks

3.1 Overview (Miah and Cuthbert)

The paper presented by Miah and Cuthbert [35] gives an excellent overview on charging schemes for ATM networks. The authors evaluate different usage-based approaches known from literature and compare them to a new charging method developed in the context of the ACTS project CANCAN.

The following design criteria for charging schemes are defined to enable comparison:

- **Clarity**: What does the user have to know about the network, the traffic and the charging method?
- **Accountability**: "Can the user's actual usage of the service be traced in response to an audit query?"
- **Predictability**: Is the tariff already known on call setup or does it depend on statistical parameters?
- **Flexibility**: Is an extension or modification of the charging scheme possible?
- **Practicality**: What implementation complexity is required for the charging scheme? Does it depend on CAC, UPC, or measurements?
- **Control**: Is the charging scheme appropriate to control network operation, e. g. to avoid congestion?
- **Choice**: Is it possible for the users to choose between different QoS levels corresponding with different prices?

3.2 Effective Bandwidth Approach (Kelly, Courcoubetis et al.)

This charging method is presented in several papers, each emphasizing a different aspect, by Kelly, Songhurst, Courcoubetis, Siris, Stamoulis, and Weber [8, 9, 10, 11, 12, 21, 22, 23, 45]. Many of the papers are based on research activities performed in the ACTS project CA$hMAN.

The basic idea is the following. Each connection is charged with a price a per until time, a price per bit b and an optional connection price c. The parameters a and b depend on the traffic characteristics negotiated during call setup following the concept of effective bandwidth. If a user declares a mean rate m for the new connection then a and b are chosen such that the straight line $a + b \cdot x$ defines the tangent of the (expected) effective bandwidth function (obtained from large deviations theory) at the point $x = m$ (Fig. 3.1). If the user actually sends traffic with a mean rate M different from m he is punished with a charge higher than that he would have to pay when he had declared the mean rate to be M. The scheme is therefore a compromise between charging for reserved resources and charging for actually needed resources.

Following drawbacks can be identified for this charging scheme [35]:

- The user has to declare the mean rate. So users who are not able to predict the mean rate of their application are punished.

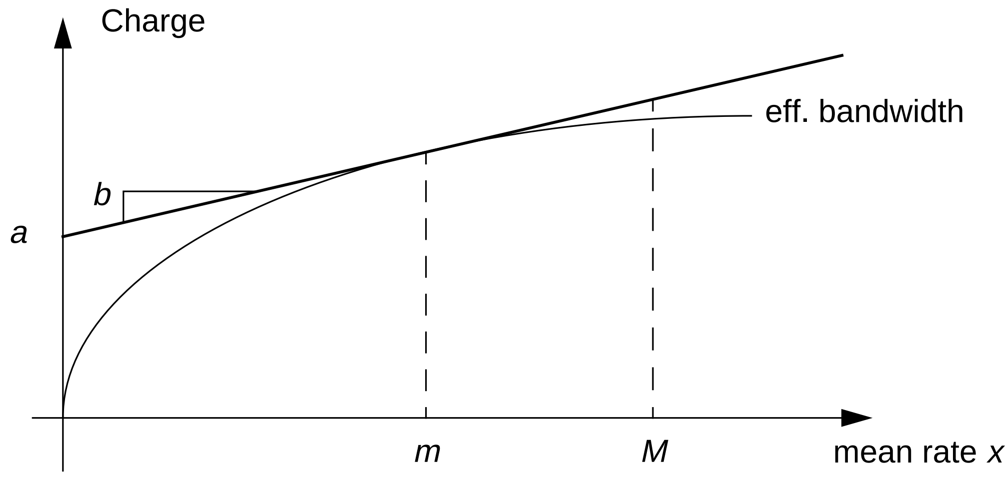

Figure 3.1: Specification of charging parameters in Kelly‘s approach

- The network has to measure the mean rate M for each connection.
- A user who is not familiar with the effective bandwidth concept will not understand the reason and the extent of the punishment.
- The CAC function has to rely on declared parameters. Therefore the UPC function will discard non-conforming cells. This may lead to a twofold punishment which will be difficult to explain to the users.

In [9] Courcoubetis et al. present an approach which enhances this model by considering different service categories and different QoS requirements. The authors, however, don‘t go into detail and so there still remain doubts whether an integration of this concept into a network will be feasible and acceptable for the users.

3.3 Effective Bandwidth Approach (Lindberger)

In [26] Lindberger proposes a simple charging method where for each connection the charge C is defined by the formula

$$C = K_{L,T} \cdot d \cdot t \tag{3.1}$$

using the following definitions:

- $K_{L,T}$: factor depending on distance and time of day
- d: a-priori estimation of effective bandwidth of the connection
- t: duration of the connection

The main problem with Lindberger‘s approach is the calculation of the effective bandwidth which is based on the bufferless multiplexer model. The rate distribution of the source may either be determined by measurement or given by the user. The first alternative is critical with respect to CAC which usually relies on the effective bandwidths being fixed during a connection. Letting the user specify the rate distribution on the other hand is too complex and has to be monitored

3.4 Mean Bandwidth Approach (Botvich)

The charging method presented by Botvich in [2] is a result of the ACTS project CANCAN. The charge for a connection is determined as follows:

$$C = \alpha \cdot K \cdot m \cdot t \tag{3.2}$$

The symbols in the formula have the following meanings:

- t: connection duration
- m: mean rate of the connection
- α: constant which may depend on time of day
- K: factor trying to predict the ratio d/m where d denotes the effective bandwidth

K is calculated by evaluating the last N released connections for which m_i, d_i, and t_i are known:

$$K = \frac{\sum_{i=1}^{N} t_i \cdot d_i}{\sum_{i=1}^{N} t_i \cdot m_i} \tag{3.3}$$

The effective bandwidths d_i are obtained from large deviations theory using a measurement of the rate distributions.

A drawback of this scheme is that the current value K has to be transmitted to the users before connection setup in order to fulfil the clarity and predictability criteria mentioned in Section 3.1. This is even more serious as different values K have to be provided for different traffic types. The choice of N and the definition of what is a traffic type are open issues which have a strong impact on both effectiveness and implementation complexity of the scheme. Furthermore, it is doubtful whether users want their connection to be charged depending on the connection parameters of previous connections of different users. This reduces the advantage that the user only has to declare the mean rate of the connection.

3.5 Design Rate Approach (Griffiths, Miah and Cuthbert)

The new charging method proposed in [35] is based on the design rate scheme presented by Griffiths in [17]. In this scheme the charge C of a connection is calculated as follows:

$$C = \alpha \cdot D \cdot t \tag{3.4}$$

where t denotes the connection duration and α is a constant factor. D represents the so-called design rate which can be specified by the user. The design rate corresponds to the rate of a traffic shaper dedicated to the connection. The shaper may be located in the end device or at the network ingress. For the case that the shaper is not under the control of the network provider the traffic leaving the shaper has to be monitored by a policer.

In [35] Miah and Cuthbert extend this concept by the support of different traffic categories (CBR/VBR, ABR, UBR) and different levels of QoS (high, medium, low). For CBR/VBR traffic the charge of a connection is given by

$$C_{CBR} = c \cdot D \cdot t \tag{3.5}$$

where c is a cost function mainly dependent on the QoS level, the time of day, and the distance between the connection end points. For UBR traffic the charge is determined by the traffic volume V independent of the connection duration:

$$C_{UBR} = c \cdot V \tag{3.6}$$

The charging function for ABR is a mixture of that for CBR/VBR and that for UBR.

The distinction of different QoS levels has to be supported by an architecture using a scheduling mechanism that is able to divide the link rate into dedicated rate shares for the different levels. The question of how to dimension the scheduler as well as the shapers is widely left open.

4 Conclusions

In the current Internet mainly flat rate pricing is used to charge users. This kind of pricing has the tremendous advantage of simplicity because no accounting is required and the effort for billing can be limited to a minimum.

However, mainly due to the following reasons usage-sensitive pricing is attractive:

- A discrimination of light users is avoided as each user pays for the amount of data sent.
- Usage-sensitive pricing helps to reduce congestion if higher per volume prices are introduced in the peak hours.
- Usage-sensitive pricing can be used as a means to fulfil different QoS requirements imposed by different types of services, applications, and users.

For the realization of the last point in the list two basic concepts are currently discussed. The first concept can be interpreted as an enhancement of the best-effort service by introducing different priority levels. The integration into the existing protocol world can easily be managed by using the TOS field in the IP header. The drawback of this approach is that no guarantees about any QoS measures like loss or delay can be given to the users. Priorities only provide a relative differentiation. The approaches presented in Section 2 widely belong to this category. Among them the Paris Metro Pricing seems to be best-reasoned with respect to feasibility.

The alternative concept is to use explicit resource reservation with guarantees that the QoS requirements specified by the user are maintained. This can be achieved by using ATM as the basic technology or integrating new reservation protocols like RSVP into the Internet. QoS guarantees may be a desirable feature for some users but the realization of that concept requires CAC. The design and management of CAC in a heterogeneous multi-service network is a complex task. Furthermore, there is always the problem whether the price of a connection should be defined according to the traffic contract or whether also the actual resource usage should be considered. A compromise solution is given in the papers from the context of the CA$hMAN project. However, the pricing scheme described there leads to rather complex mathematical models which can hardly be understood by a normal user. A convincing solution for the issue of different QoS requirements is also not available.

The shortcomings of a concept based on reservation and QoS guarantees are obvious. The models proposing some kind of priority pricing on the other hand have the advantage that they can be more easily integrated into the current structure of the Internet, thus offering a migration path for providers. Issues like the realization of a “receiver pays” scheme or the pricing of multicast services still remain for further study. Moreover, architectures of network elements have to be found which support multiple priority traffic in an appropriate way.

References

[1] J. P. Bailey, L. W. McKnight, H. S. Sharifi: "Critical Business Decisions for Internet Services", *Proceedings of the 1998 Information Resources Management Association International Conference*, Boston, May 1998, http://itel/mit/edu/itel/publications.html.

[2] D. Botvich, Y. Chen, T. Curran, B. Kerswell, J. McGibney, D. Morris: "On Charging for Internet Services Provided over an ATM Network", *Proceedings of the IEEE ATM '97 Workshop*, Lisboa, Portgal, pp. 669-679, May 1997.

[3] N. Brownlee: "Internet Pricing in Practice", in *Internet Economics*, eds. L. W. McKnight, J. P. Bailey, pp 77-90, 1995, http://www.press.umich.edu/jep/econ-TOC.html.

[4] S. Chatterjee, C. S. Thachenkary, J. L. Katz: "Modeling the economic impacts of broadband residential services", *Computer Networks and ISDN Systems*, Vol. 30, No. 14, August 1998.

[5] D. D. Clark: "Internet Cost Allocation and Pricing", in *Internet Economics*, eds. L. W. McKnight, J. P. Bailey, pp 215-278, 1995, http://www.press.umich.edu/jep/econ-TOC.html.

[6] R. Cocchi, D. Estrin, S. Shenker, L. Zhang: "A Study of Priority Pricing in Multiple Service Class Networks", *Proceedings of ACM SIGCOMM '91*, Zurich, September 1991, ftp://parcftp.xerox.com/pub/net-research/pricing1.ps.Z.

[7] R. Cocchi, S. Shenker, D. Estrin, L. Zhang: "Pricing in Computer Networks: Motivation, Formulation, and Example", *IEEE/ACM Transactions on Networking*, vol. 1 , no. 6, pp. 614-627, 1993.

[8] C. Courcoubetis, V. A. Siris, G. D. Stamoulis: "Comparing Usage-Based Pricing Schemes for Broadband Networks", *Proceedings of the IEE Colloquium "Charging for ATN - The Reality Arrives"*, 1997, http://www.ics.forth.gr/proj/race/publications/index.html.

[9] C. Courcoubetis, V. A. Siris, G. D. Stamoulis: "Charging Issues for Services in Broadband Networks", *Proceedings of the OECD Conference on the Economics of the Information Society*, Istanbul, December 1995, http://www.ics.forth.gr/proj/race/publications/index.html.

[10] C. Courcoubetis, V. A. Siris: "An Evaluation of Pricing Schemes that are based on Effective Usage", *Proceedings of the IEEE International Conference on Communications (ICC '98)*, 1998, http://www.ics.forth.gr/proj/race/publications/index.html.

[11] C. Courcoubetis, F. P. Kelly, R. Weber: "Measurement-Based Charging in Communications Networks", *Technical Report*, 1997, http://www.ics.forth.gr/proj/race/publications/index.html.

[12] C. Courcoubetis, F. P. Kelly, V. A. Siris, R. Weber: “A study of simple usage-based charging schemes for broadband networks”, *Proceedings of the IFIP TC6 / WG6.2 Fourth International Conference on Broadband Communications (BC ‘98)*, Stuttgart, Germany, pp 209-221, 1998, http://www.ics.forth.gr/proj/race/publications/index.html.

[13] G. de Veciana, R. Baldick: “Resource allocation in multi-service networks via pricing: statistical multiplexing”, *Computer Networks and ISDN Systems*, vol. 30 , no. 9-10, pp. 951-962, 1998.

[14] R. Edell, N. McKeown, P. P. Varaiya: “Billing Users and Pricing for TCP”, *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7, pp. 1162-1175, 1995.

[15] D. Ferrari, L. Delgrossi: “Charging For QoS”, *Proceedings of the 6th IEEE/IFIP Workshop on Quality of Service (IWQoS ‘98)*, Napa, California, USA, May 1998, http://www-ece.rice.edu/conf/iwqos98.

[16] P. C. Fishburn, A. M. Odlyzko: “Dynamic Behavior of Differentiated Pricing and Quality of Service Options for the Internet”, *Proceedings of the 1st International Conference on Information and Computation Economics (ICE 98)*, 1998, pp. 128-139, http://www.research.att.com/~amo/doc/networks.html.

[17] J. Griffiths: “Poisson Envelope Charging for ATM: Use of Multidimensional Markov Chain Animation”, *2nd IFIP Workshop on Traffic Management, Synthesis of ATM Networks*, Montreal, Septemeber 1995.

[18] A. Gupta, D. O. Stahl, A. B. Whinston: “Priority Pricing of Integrated Services Networks”, in *Internet Economics*, eds. L. W. McKnight, J. P. Bailey, pp 323-352, 1995, http://www.press.umich.edu/jep/econTOC.html.

[19] D. Hazlett: “An Interim Solution to Internet Congestion”, *Whitman College Internal Report*, Walla Walla, WA, http://marcus.whitman.edu/pub/Hazlett/congestion.ps.

[20] S. Kalyanaraman, T. Ravichandran, R. Norsworthy: “Dynamic Capacity Contracting: A Framework for Pricing the Differentiated Services Internet”, submitted to *1st International Conference on Information and Computation Economics*, 1998, http://networks.ecse.rpi.edu/~shivkuma/tm-papers.html.

[21] F. P. Kelly: “Charging and Accounting for Bursty Connections”, in *Internet Economics*, eds. L. W. McKnight, J. P. Bailey, pp 253-278, 1995, http://www.press.umich.edu/jep/econTOC.html.

[22] F. P. Kelly: “Charging and Rate Control for Elastic Traffic”, *European Transactions on Telecommunications (ETT)*, vol. 8, no. 1, pp. 33-37, 1997, http://www.statslab.cam.ac.uk/~frank/PAPERS.

[23] F. P. Kelly: “Tariffs and effective bandwidths in multiservice networks”, *Proceedings of the 14th Internation Teletraffic Congress (ITC 14)*, France, pp 401-410, 1994.

[24] K. Kilkki: “Simple Integrated Media Access - a Novel Service Concept for Internet”, *Proceedings of the ATM Traffic Symposium*, Mykonos, Greece, Sep. 1997.

[25] B. C. Klopfenstein: “Internet Economics: An Annotated Bibliography”, 1995, http://www.bgsu.edu/departments/tcom/anota.htm.

[26] K. Lindberger: “Cost Based Charging Principles in ATM Networks”, *Proceedings of the 15th International Teletraffic Congress (ITC 15),* Washington D.C., USA, pp. 771-780, June 1997.

[27] J. K. MacKie-Mason, L. Murphy, J. Murphy: “Pricing Congestible Networks”, *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7, pp. 1141-1149, 1995, http://www-personal.umich.edu/~jmm/papers.html.

[28] J. K. MacKie-Mason, H. R. Varian: “Economic FAQs About the Internet”, in *Internet Economics*, eds. L. W. McKnight, J. P. Bailey, pp 27-62, 1995, http://www.press.umich.edu/jep/econTOC.html.

[29] J. K. MacKie-Mason, L. Murphy, J. Murphy: “Responsive Pricing in the Internet”, in *Internet Economics*, eds. L. W. McKnight, J. P. Bailey, pp 279-304, 1995, http://www.press.umich.edu/jep/econTOC.html.

[30] J. K. MacKie-Mason, H. R. Varian: “Pricing the Internet”, in *Public Access to the Internet*, pp. 269-314, 1995, http://www-personal.umich.edu/~jmm/papers.html.

[31] J. K. MacKie-Mason: “A Smart Market for Resource Reservation in a Multiple Quality of Service Information Network”, *Internal Report*, University of Michigan and NBER, Ann Arbor, USA, September 1997, http://www-personal.umich.edu/~jmm/papers.html.

[32] L. W. McKnight, J. P. Bailey: *Internet Economics*, The MIT Press, Cambridge, Mass., 1997.

[33] L. W. McKnight, J. P. Bailey: “An Introduction to Internet Economics”, in *Internet Economics*, eds. L. W. McKnight, J. P. Bailey, pp 3-26, 1995, http://www.press.umich.edu/jep/econTOC.html.

[34] L. W. McKnight, B. Leida: “Internet Telephony: Costs, Pricing, and Policy”, *Proceedings of the 25th Annual Telecommunications Policy Research Conference*, Alexandria, VA, September, 1997, http://itel/mit/edu/itel/publications.html.

[35] B. Miah, L. Cuthbert: “Charging and Billing in ATM Networks”, *Proceedings of the 6th IFIP Workshop on Performance Modelling and Evaluation of ATM Networks*, Ilkley, West Yorkshire, UK, pp. 92/1-92/9, July 1998.

[36] L. Murphy, J. Murphy: “Feedback and Pricing in ATM networks”, *Proceedings of the 3rd IFIP TC6 Workshop on Performance Modelling and Evaluation of ATM Networks,* Ilkley, UK, July 1995, http://www.eeng.dcu.ie./~murphyj/publ/publ.html.

[37] A. Odlyzko: “A modest proposal for preventing Internet congestion”, *AT&T Labs Internal Report*, Florham Park, NJ, September 1997, http://www.research.att.com/~amo/doc/complete.html.

[38] A. Odlyzko: “The economics of the Internet: Utility, utilization, pricing, and Quality of Service”, *AT&T Labs Internal Report*, Florham Park, NJ, July 1998, http://www.research.att.com/~amo/doc/networks.html.

[39] A. Odlyzko: “The Internet and other networks: Utilization rates and their implications”, *AT&T Labs Internal Report*, Florham Park, NJ, September 1998, http://www.research.att.com/~amo/doc/networks.html.

[40] A. Orda, N. Shimkin: “Incentive Pricing in Multi-Class Communication Networks”, *Proceedings of IEEE Infocom ‘97*, pp. 659-666, 1997.

[41] J. W. Roberts: “Quality of Service Guarantees and Charging in Multiservice Networks”, *IEICE Transactions on Communications*, vol. E81-B , no. 5, pp. 824-831, 1998.

[42] H. Saito: “Resource Management and Charging in ATM Networks”, *Computer Networks and ISDN Systems*, vol. 28, no. 5, pp 641-644, 1996.

[43] S. Shenker: “Fundamental Design Issues for the Future Internet”, *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7, 1995.

[44] S. Shenker: “Service Models and Pricing Policies for an Integrated Services Internet”,, ftp://parcftp.xerox.com/pub/net-research/policy.ps.Z.

[45] D. Songhurst, F. Kelly: “Charging Schemes for Multiservice Networks”, *Proceedings of the 15th International Teletraffic Congress (ITC 15)*, Washington D.C., USA, pp. 781-790, June 1997, http://www.statslab.cam.ac.uk/~frank/PAPERS.

[46] B. Stiller, G. Fankhauser, B. Plattner: “Charging of Multimedia Flows in an Integrated Services Network”, *Proceedings of the 8th International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV 98)*, Cambridge, UK, July 1998, pp. 189-192, http://www.cl.cam.ac.uk/Research/SRG/nossdav98/#tech.