# Charging and Accounting in High-speed Networks

Burkhard Stiller[1], Peter Reichl[2], Jan Gerke[1], Hasan[1], Placi Flury[1]

*[1] Computer Engineering and Networks Laboratory TIK,*
*Swiss Federal Institute of Technology, ETH Zürich, Switzerland*
*[2] Forschungszentrum Telekommunikation Wien FTW, Vienna, Austria*

**Abstract**

The commercialization of the Internet and the requirement to offer various different network services for emerging applications has driven the development of appropriate Internet technology, such as the Differentiated Services Architecture. However, these developments focussed on the technological point of view only, neglecting the need to provide economic incentives to users and customers to chose the "right" service class. Once accepted that charging Internet services is an economically valid approach, suitable charging technology is required. Therefore, this paper motivates and introduces a suitable Internet Charging System. In addition, the development of incentive compatible pricing models for multi-service networks is the corresponding economic question. The problem of today's flat rates shows, that high-speed networks are priced simply and customers like them, however, Internet Service Providers tend to show a severe economic and technical instability in the market. Therefore, a better than flat fee approach termed Cumulus Pricing Scheme is introduced, solving the feasibility problem of pricing Internet services.

## 1 Introduction

As current trends determine, the Internet may become a ubiquitious, global network for communications. Internet services will accommodate communications between various locations and users. Basic best-effort services as well as advanced and guaranteed services are required to support the broad range of today's multimedia and e-commerce applications. However, while enhanced service concepts emerge on the technical side, a wide-range and professional provision of Internet services is somewhat restricted due to lacking mechanisms for the efficient and usage-based charging of integrated Internet services. Prerequisites to perform such a usage-based charging are manyfold. On one hand, the technology in terms of metering and accounting devices is necessary to obtain hooks for the measurement of service usage and its detailed recording. While the Network Traffic Meter (NeTraMet) [5] represents an example for metering devices, Internet Protocol Detail Records (IPDR) [9] define an upcoming data unit format, which deals with packet-based

[1] E-Mail: `[stiller|gerke|hasan|flury]@tik.ee.ethz.ch`

[2] E-Mail: `reichl@ftw.at`

Internet protocol details. On the other hand, pricing models, charge calculation procedures, and customer data are required to calculate a usage-sensitive price for a utilized service, which has to be paid for by the customer utilizing this service. The term charging will be used at this stage for the overall process of identifying resource usage until an invoice is presented to the customer. Accounting covers technical as well as economic areas, however, the technical accounting is the bottleneck to be considered for high-speed networks.

Charging and accounting are those functions required in a commercially operated network, which allow for the support of usage-sensitive service management. While charging traditionally used to be performed on a flat fee basis, a significant number of Internet Service Providers (ISP) went bankrupt in recent months due to this pricing scheme and its lack of incentive compatibility. Therefore, to close the gap between today's Internet communication demands and the need to gain revenue for ISPs, economic incentives and technical support are required in an integrated manner. In this sense, pricing is considered the most important management functionality future commercial networks need to offer, following actual market forces. Since the Internet is on the move to provide differentiated services, for the backbone based on the Differentiated Services Architecture (DiffServ) [2], suitable and scalable management mechanisms are required.

For that reason, this overview discusses key problems in the area of Internet charging, introduces a charging terminology for the Internet, and sketches the DiffServ architecture and its Quality-of-Service (QoS) model. In particular, it provides insights in the dilemma of time-scales, which arises in case of the definition of new and, at the same time, technically feasible pricing models for Internet communication services. This leads to the design of a framework for pricing models and their application in a segmented ISP market. It is the only framework known so far defining a clear relation between different time-scales of accounting periods, measurement periods, and charging periods. Based on this new view of pricing, considered as an extension to traditional management information, the Cumulus Pricing Scheme (CPS) proposed targets particularly at DiffServ technology. Prices in this scheme are based on flat fees and, hence, are predictable and transparent. CPS solves the "Internet Feasibility Problem", the solution of a trade-off between customer demands, ISP economic optimizations, and ISP's technical efficiency. In support, a generic and modular Internet Charging System (ICS) has been designed, which offers a communication service-independent architecture and integrates price-driven, economically-controlled network management functions. The ICS may perform transport, service, and content charging, it supports accounting tasks according to transport and multi-service definitions, and is able to integrate different levels of security for charging information.

## 2 Terminology and Charging Multi-service Networks

Charging terminology, especially for the Internet, is still in flux and requires a careful definition before applying it to existing systems or the one developed. Therefore, the unique terminology will be used as follows [22], while describing at the same time a sequential step of procedures to be performed for preparing a customer's bill for his service use. Charging is used within this work as an overall term, depicting at a higher layer of abstraction all tasks required to calculate the finalized content of a billing record. Sometimes in the literature, the term billing is utilized instead. On the technical side metering determines the capture of usage of resources within hosts or routers, applications, calls, or any type of connections on a technical level, including QoS, management, and networking parameters. Accounting defines the summarization of these information in terms of accounting records and in relation to a customer's service utilization. The mediation may be interleaved and is intended to filter, aggregate, and correlate raw technical data coming from a meter. Mediation transforms these data into a form which can be used for storing and further processing in the accounting task. Concerning the economic input variables, pricing covers the specification and setting of prices for goods, specifically networking resources and services. This process may combine technical considerations, *e.g.*, resource consumption, and economical ones, *e.g.*, applying tariffing theory or marketing. Prices may be calculated on a cost/profit base or on the current market situation. A tariff defines the algorithm used to determine a charge for a service usage and may contain, *e.g.*, discount strategies, rebate schemes, or marketing information. It is applied in the charge calculation for a given customer for the service he utilizes. Therefore, the charge calculation covers the complete calculation of a price for a given accounting record and its consolidation into a charging record, while mapping technical values into monetary units. Billing defines the collection of charging records, summarizing their monetary content, and delivering a bill or invoice to a user, including an optional list of detailed charges and service information.

## 2.1 Charging Multi-service Networks

Over the last couple of years it has been argued extensively that pricing in general and usage-based pricing in particular imposes a high overhead on telecommunication systems [15], [20]. However, due to incentive compatibility reasons, a form of usage-based pricing for Internet services is essential in a commercialized environment. This becomes even more evident by considering a multi-service Internet, where QoS-enhanced services are offered at the same time with best-effort services. Obviously relevant, Internet services utilizing underlying resources (such as satellites, frequencies, cables, routers/switches, and most notably operating personnel) at a larger degree than others have to be paid for at a higher price [7]. These mentioned resources are and will remain scarce and costly, even though the world-wide Internet access may become a commodity. This is mainly due to areas of high investments for installing the infrastructure (trans-atlantic or trans-pacific undersea cables) and areas of generally low demand, but overutilized existing equipment (such in less populated or over-populated regions of the world). Since in any technically operated system the waste of resources is unacceptable, unfairness among users has to be tackled by a clearly communicated policy. The loss of revenues for ISPs is an undesirable situation, incentive-compatible charging approaches for multi-service networks are a must.

A flat fee pricing model for those services shows economic drawbacks of not being incentive-compatible [6], [8], [13], and [20]. Therefore, a per-flow billing system for TCP (Transmission Control Protocol) flows had been developed some time ago [11]. Going beyond this work, more advanced per-flow charging and accounting approaches based on reservations have been tackled in [10], [12], and [17] to cope with the Integrated Services Architecture [3]. Multi-service packet networks are proposed by adding classification and scheduling to routers, but not policing [4]. Concerning pricing for multi-service networks, it was stated that users were not disinclined to flexible pricing models [1]. Finally, the project M3I (Market-Managed Multi-service Internet) [18] aims at the design and implementation of a next-generation system enabling Internet resource management through market forces, specifically by enabling differential charging for multiple service levels. This paper presents some results from the M3I project.

### 2.2 Charging Differentiated Services

The DiffServ architecture [2] for Internet backbones determines a suitable technology in support of multiple service classes and QoS. Avoiding a single per-flow handling, aggregates of application flows are mapped onto an appropriate service class and every packet is marked accordingly. However, as soon as these distinct service classes exist, an incentive for using the "right" and best-suited class is required for customers. One of such incentives will be the price to be paid for the utilization of a particular communication service. In addition, since the exchange of IP packets will not happen only between single customers and ISPs, it will also be required between ISPs as well as ISPs and Customer Premises Networks.

In all of these cases, the exchange of data and traffic based on service classes is supported by Service Level Agreements (SLA) [23], which form the technical and legal container for inter-provider and provider-customer relationships. Since the level of QoS can vary, SLAs are to be negotiable [21]. Due to the fact that the Internet traffic and its volume is changing rapidly in much shorter time-scales compared to the traditional telecommunications sector, SLA negotiations become essential. Assuming that the solution for a negotiation procedure for SLA exists in technical parameter dimensions, prices and their economic incentives are to be handled as well. Therefore, the investigation of the problem of differentiated services charging is required.

## 3 Framework for Pricing

This work on a framework for pricing Internet services proposes a two-part solution. On one hand, suitable charging technology in terms of an Internet Charging System (ICS) is presented and discussed briefly with respect to its interactions between the customer, the ISP, and the network technology [22]. The developed prototype manages tariffs and accounted for data to be applied by an ICS-internal data base. On the other hand, a DiffServ-capable charging scheme, the Cumulus Pricing Scheme CPS, is introduced [21]. CPS provides a flat fee like user interface, but offers those incentives to users that force them to act economically serious with respect to resource usage and related service pricing. The major synergy of these two parts, technology and pricing model, becomes visible in the pricing model's applicability to distinct pricing thresholds for identifying different service class over or under-utilization independently. Therefore, based on concrete numerical values of a networking scenario, pricing acts for the provider as a means to perform traffic control economically.

### 3.1 Internet Charging System

A charging system for the Internet provides the technology, which maintains data records containing information on service usage and which provides calculation routines to compute service costs to be paid for by the customer. An ICS has been developed by the authors within the M3I project [18], [22], where its core performs major functions of accounting, charge calculation, session control, and contract control. This core is referred to as the Charge Calculation and Accounting System (ICCAS).

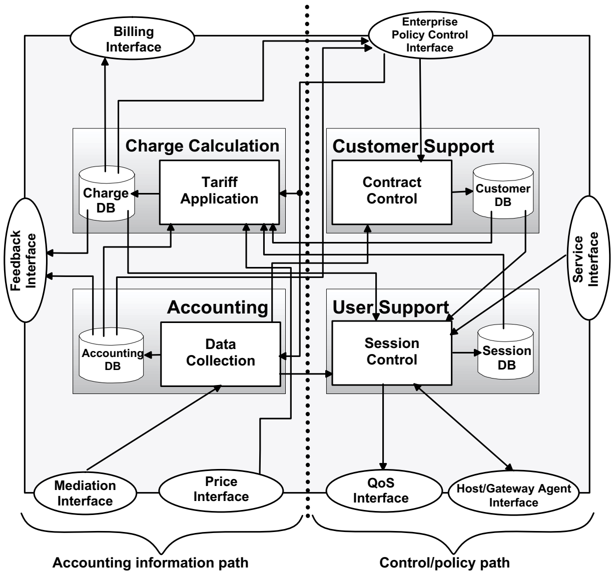


Figure 1: ICCAS Architecture in Detail

The ICCAS architecture and its basic components are presented in Figure 1. ICCAS-internal entities consist of a charge calculation, an accounting, a customer support, and a user support component. The separation of the ICCAS into these components increases the required degree of flexibility, since they can be distributed and replicated physically. On one hand, the Accounting Information Path deals with the flow of pure charging-relevant data. On the other hand, the Control Policy Path is used to manage and configure the ICCAS and those entities involved with the processing of charging data. These two paths differ mainly in the order and direction they process data. The accounting component receives all metered and mediated usage data and is responsible for storing it. It has to provide these stored data to other ICCAS components and interfaces for further processing, feedback, or statistic evaluation. The charge calculation component processes the accounted for usage data. It calculates appropriate charges for resource usage applying a tariff, communicated by the pricing component. To be able to determine the charges fully, it needs input from the user support, *e.g.*, user identifications. Since a customer is not the same as a user, *e.g.*, one customer might pay bills of several users, a customer defines the role which negotiates a contract with the ISP. The content of this contract, *e.g.*, number of users covered by the contract and their names and accounts as well as SLAs, are managed within the customer support. While the customer support component is responsible for keeping all contract information the user support component is responsible for ensuring that those contracts are kept. On one hand, this means that user requests are blocked, if they are not covered by the contract of the customer, who pays for the user. On the other hand, it has to be ensured that a service requested by a user is delivered to him, if the contract allows it.

### 3.2 Timing Dimensions

The major impact factor for a precise investigation of pricing models for the Internet is given by time-scales. Based on traditional time-scales for network management [14], the extension on the atomic time-scale in the range of milliseconds/seconds allows for the definition of four different time-scales as being relevant for Internet tariff schemes. They have been introduced in [19] and describes additionally requirements and existing mechanisms. The applicability of Internet pricing schemes depends on the balance between these requirements outlined below.

The requirements for a best-suited pricing model are driven by technical and economic aspects of ISPs and customers' perspectives on pricing schemes. On one hand, as depicted at the bottom of Figure 2, requirements of ISPs with respect to their technical infrastructure (technical feasibility), mainly metering and accounting technology, would prefer a precise usage-based approach for a technical accounting in the time-scale of milliseconds/seconds. Concerning ISPs and their economic optimization, the economic efficiency of a service provisioned could be achieved in the time-scale of seconds to days. Finally, customers expect from a pricing model a transparency and predictability in the time-scale of days to years. *E.g.*, flat rate pricing is excellent with respect to customers and ISP accounting technology, but does not support the ISP's economic efficiency at all, whereas Vickrey Auctions are economically efficient, but may yield unstable charges as well as complex technical problems.

On the other hand, existing mechanisms in support of current Internet pricing are depicted at the top of Figure 2. Firstly, current metering technology for ISPs shows a handleable efficiency in terms of data volume to be stored in the time-scale of a metering sample per minute or better hour, a per-second sample is still a major technological challenge. Certainly, these scales depend on the bandwidth of the underlying network technology. High-speed networks transmitting data in the bandwidth range of many gigabit per second are currently unable to meter fine granular data traffic. This fact probably will remain for quite some time, since additional large-scale investments in metering technology define mainly an economic problem for ISPs. Secondly, suitable economic feedback signals from customers to ISPs, such as (non)-use of service, service utilization statistics, or regular payments of the accumulated service fees, exist in the range of hours to weeks, but certainly not below these time-scales. Thirdly, customers prefer contracts of monthly or yearly durations.

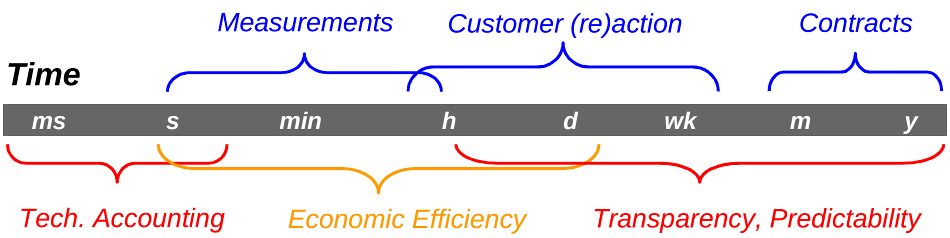


Figure 2: Time-scales of Requirements (Bottom) and Existing Mechanisms (Top)

A closer analysis of requirements and existing mechanisms together with practical experiences reveals that not all of three requirements are of equal importance. Any scheme that does not fulfill the criterion of best technical feasibility has no opportunity of being implemented for practical purposes. In this sense, various pricing schemes will have different trade-offs between these requirements. However, the technical feasibility represents the hard criterion. This fact has been termed the “Internet Feasibility Problem” for pricing [19].

### 3.3 The Cumulus Pricing Scheme CPS

In solution of this Internet Feasibility Problem, CPS has been developed in [19], [22]. It is basically a flat rate scheme, but rates may vary over long time-scales. CPS provides a feedback mechanism to bring market forces into play, where this feedback is not an immediate one, but requires the accumulation of discrete “flags” according to user behavior. Finally, CPS allows a huge flexibility in terms of the technical prerequisites for metering and accounting mechanisms.

The major simplicity feature of this scheme is the combination of an initial contract between the customer and the ISP, which contains information on expected usage patterns. It is complemented by a feedback mechanism between customer and ISP that interacts with the customer behavior on different time-scales. ISP measurements take place over a shorter time-scale (*e.g.*, once every hour) and allow sufficient evidence on user behavior on a medium time-scale. This evidence is expressed in terms of discrete Cumulus Points (CPs), yet not triggering some sort of reaction by themselves right away, but only as a result of their accumulation over a long time-scale. A customer and ISP are supposed to agree on a contract specifying (C1) the expected user requirements in terms of, *e.g.*, bandwidth or delay, (C2) a flat rate to be paid for this type of service over, *e.g.*, a week or month, and (C3) a contract duration, which may be extended automatically by the same duration. Following this agreement, the actual usage may not match the prediction given by the user, may be due to an incorrect statement, changing habits, or new applications in use. As soon as these discrepancies exceed a defined threshold as defined relatively or absolutely according to the C1 requirements, the user receives feedback in terms of these mentioned CPs as defined in C2. They exist as red and green economic feedback signals. A red CP indicates that the user has overused the capacities negotiated, a green one indicates the opposite, i.e. that the user might have been allowed to use more resources than actually used. The larger the discrepancy between contract and measured data, the more CPs may be assigned. CPs remain valid for a dedicated number of consecutive periods, according to C3, which have been negotiated in the initial contract for billing purposes, *e.g.*, every week or month. Only their final accumulation triggers consequences between the customer and the ISP. Hence, receiving CPs requires no immediate reaction. However, their successive accumulation over consecutive contract periods eventually may exceed the CP threshold and will have consequences for the user, depending on various different ISP policies.

A simulation model has been implemented for CPS to perform an exhaustive simulative evaluation of major CPS aspects. The following results on the assignment of CPs depending on the metering frequency are based on measured network data traffic over 240 working days between a TIK Lab Ethernet LAN and the ETH Zürich backbone. For various granularities (ranging from 2 minutes to 12 hours) the current bandwidth consumption is determined once per granularity interval, *e.g.*, once within each 2 minute interval at a randomly chosen instant. It is assumed that the customer has delivered a correct initial specification (*i.e.* mean $\mu$ and standard deviation $\sigma$) of the expected traffic, and three CP assignment threshold levels for 1, 2, or 3 CPs relatively based on the standard deviation $\sigma$ have been chosen particularly to be $1{,}3\sigma$, $2{,}4\sigma$, and $3{,}1\sigma$. For this scenario, Figure 3 (left) depicts the influence of measurement granularity, assuming that the initial specification was correct. From this and similar results it can be concluded that CPS is very stable for granularities up to 2 to 3 hours (corresponding to at least 8 to 12 suitably distributed measurements per day), whereas for larger granularities the outcome is less predictable. Figure 3 (right) shows the situation, where the customer specification is underestimating his real requirements by 10%. In this case, the reaction of CPS on all granularity levels is consistent, assigning red CPs. In every case, a suitable reaction threshold level, *e.g.*, 5 CPs, is reached at least after 3 to 5 months of continuously exceeding the specified requirements.

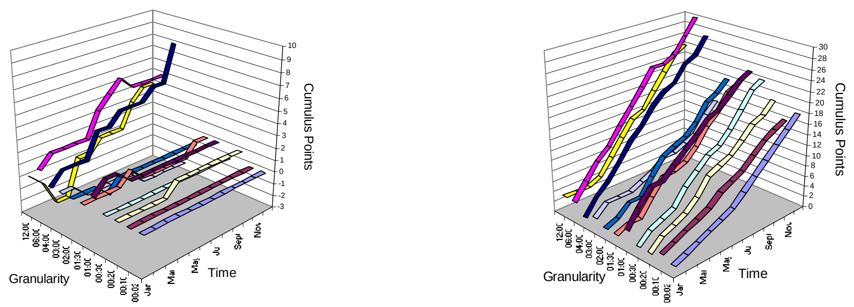


Figure 3: CP Assignment Results for Correct Traffic Estimation (left) and a 10% Underestimation

## 4 Summary and Conclusions

This overview on charging high-speed networks defined a unique terminology on charging for Internet services. Based on a brief introduction of Internet charging problems in general, the major reason for this problem was presented, mainly the missing incentive compatibility of flat rate pricing schemes for Internet service users and a large effort for usage-based accounting technology. Pricing Differentiated Services (DiffServ) for multi-service networks based on Service Level Agreements (SLA) was motivated and a pricing framework for dealing with different time-scales has been introduced. Those timing dimensions between customers and Internet Service Providers (ISP) were investigated in different time-scales with respect to their requirements and existing mechanisms in support of pricing models. The derived Internet Feasibility Problem of pricing was solved by the Cumulus Pricing Scheme (CPS), which offers a better than flat fee pricing approach for Internet services. CPS achieves a customer transparency including economic feedback and essentially it reduces the measurement overhead from ISPs.

Concluding, the necessity to support differentiated services pricing becomes even more stringent as a variety of Internet services is offered in an open market. Therefore, incentive compatible pricing schemes will replace existing flat fee pricing schemes, otherwise ISPs will reduce to bankruptcy [16]. ISPs will need to follow a technically efficient way and require an implementation of an economically efficient scheme, which is offered by CPS. Future work, which was already started recently, encompasses additional investigations of CPS with respect to its implementation in a larger network, with dedicated DiffServ service classes, and further economic stability issues.

## References


[1] J. Altmann, B. Rupp, P. P. Varaiya: *Effects of Pricing on Internet User Behavior;* Netnomics, Vol. 3, No. 1, June 2001, pp 67-84.

[2] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss. *An Architecture for Differentiated Services;* Internet Engineering Task Force, RFC 2475, December 1998.

[3] R. Braden, D. Clark, S. Shenker: *Integrated Services in the Internet Architecture: An Overview;* Internet Engineering Task Force, RFC 1633, June 1994.

[4] B. Briscoe, M. Rizzo, J. Tassel, K. Damianakis: *Lightweight Policing and Charging for Packet Networks;* 3rd IEEE Conference on Open Architectures and Network Programming (OpenArch 2000), Tel Aviv, Israel, March 26-27, 2000.

[5] N. Brownlee: *Traffic Flow Measurement: Experiences with NeTraMet;* Internet Engineering Task Force, RFC 2123, March 1997.

[6] D. Clark: *Combining Sender and Receiver Payments in the Internet*; Telecommunications Research Policy Conference, October 1996.

[7] Cisco Systems: *Service Providers New World News - Pricing for Profitability;* Packet, Cisco Systems, 3rd Quarter 2000, pp 35–39.

[8] R. Cocchi, D. Estrin, S. Shenker, L. Zhang: *Pricing in Computer Networks: Motivation, Formulation and Example;* IEEE/ACM Transactions on Networking, Vol. 1, No. 6, December 1993, pp 614–627.

[9] S.A. Cotton (edt.): *Network Data Management – Usage (NDM-U) for IP-Based Services;* IPDR Specification Version 1.1, June 2000.

[10] G. Dermler, M. Günter, T. Braun, B. Stiller: *Towards a Scalable System for Per-flow Charging in the Internet;* Applied Telecommunication Symposium, Washington D.C., U.S.A., April 17-19, 2000.

[11] R. Edell, N. McKeown, P. P. Varaiya: *Billing Users and Pricing for TCP;* IEEE Journal an Selected Areas in Communications, Vol. 13, No. 7, July 1995, pp 1162–1175.

[12] G. Fankhauser, B. Stiller, C. Vögtli, B. Plattner: *Reservation-based Charging in an Integrated Services Network;* 4th INFORMS Telecommunications Conference, Bocca Raton, Florida, U.S.A., March 1998.

[13] A. Gupta, D. O. Stahl, A. B. Whinston: *Managing the Internet as an Economic System;* CISM, University of Texas at Austin, U.S.A., July 1994.

[14] H.-G. Hegering, S. Abeck, B. Neumair: *Integrated Management of Networked Systems;* Morgan Kaufmann Publishers, San Francisco, California, U.S.A., 1999.

[15] J. MacKie-Mason, H. Varian: *Pricing Congestible Network Resources;* IEEE Journal on Selected Areas in Communications, Vol. 13, No. 7, 1995, pp 1141–1149.

[16] N. McIntosh: *Heavy Surfers Pay the Price.* The Guardian, England, U.K., August 3, 2000.

[17] M. Karsten, J. Schmitt, L. Wolf, R. Steinmetz: *An Embedded Charging Approach for RSVP;* 6th IEEE/IFIP International Workshop on QoS (IWQoS’98), Napa, California, U.S.A., May 18-20, 1998.

[18] M3I: *Market Managed Multi-service Internet;* 5th Framework EU Project No. IST-1999-11429, available at the URL: http://www.m3i.org, March 2001.

[19] P. Reichl, P. Flury, J. Gerke, B. Stiller: *How to Overcome the Feasibility Problem for Tariffing Internet Services: The Cumulus Pricing Scheme;* IEEE International Conference on Communications (ICC 2001), Helsinki, Finland, June 11-15, 2001.

[20] S. Shenker, D. Clark, D. Estrin, S. Herzog: *Pricing in Computer Networks: Reshaping the Research Agenda;* ACM Computer Communication Review, Vol. 26, No. 2, April 1996, pp 19–43.

[21] B. Stiller, J. Gerke, P. Reichl, P. Flury: *Management of Differentiated Service Usage by the Cumulus Pricing Scheme and a Generic Internet Charging System;* 7th IEEE/IFIP Symposium on Integrated Network Management (IM 2001), Seattle, Washington, U.S.A., May 14-17, 2001.

[22] B. Stiller, P. Reichl, J. Gerke, P. Flury: *A Generic and Modular Internet Charging System for the Cumulus Pricing Scheme;* to appear: Journal on Network and Systems Management, September 2001.

[23] B. Teitelbaum, P. Chimento: *QBone Bandwidth Broker Work Group.* available at the URL: http://qbone.ctit.utwente.nl/BBroker/, March 2001.