

Pricing and Cost Recovery for Internet Services: Practical Review, Classification, and Application of Relevant Models

Burkhard Stiller¹, Peter Reichl¹, Simon Leinen²

¹Computer Engineering and Networks Laboratory TIK, ETH Zürich,

Gloriastrasse 35, CH – 8092 Zürich, Switzerland, E-Mail: [stiller | reichl]@tik.ee.ethz.ch

²SWITCH, Zürich, Limmatquai 138, CH – 8001 Zürich, Switzerland, E-Mail: leinen@switch.ch

Abstract

Suitable pricing models for Internet services represent one of the main prerequisites for a successfully running implementation of a charging and accounting system. This paper introduces general aspects influencing the choice of a pricing model in practical situations and presents a survey as well as a classification of relevant and advanced approaches to be found in the scientific literature. First performance results on charging extensions within the Internet are presented, which are completed by a set of market price simulations for dynamic pricing models within the same implementation environment. Based on cost model investigations some detailed insights into price and cost issues from an Internet Service Provider's (ISP) point of view are given. Moreover, current challenges as well as problems are discussed in a practical context as investigated in the Swiss National Science Foundation project Charging and Accounting Technology for the Internet (CATI).

Keywords: Internet Pricing, Auction Pricing, Cost Recovery, Peering Agreements.

1 Introduction

It is expected that many Electronic Commerce applications will make use of the Internet as a transport infrastructure. In contrast to more traditional communication networks such as the Public Switched Telephone Network (PSTN), the current public Internet lacks well-defined pricing models and cost recovery schemes, both between Internet Service Providers (ISPs) and their clients and between interconnected ISPs.

As Electronic Commerce becomes more widespread, businesses will increasingly recognize the impact of the quality of the Internet infrastructure for revenue. This will spur a competition for improved Internet service between connected businesses, which will in turn lead to differentiated Internet service offerings from ISPs. As the Internet will move away from a pure "best-effort" service to a more differentiated service network, new schemes for pricing and Service Level Agreements (SLA) will necessarily evolve. A major goal of this work is to anticipate and investigate such schemes.

As the variety of Internet services, such as native IP (Internet Protocol) service, Mbone service, WWW (World Wide Web) services, or FTP (File Transfer Protocol) service, show extremely different traffic characteristics, sensible pricing models for integrated and differentiated services and for multicast traffic [4] in the Internet are essential. Therefore, the development of appropriate SLAs, pricing models, and suggested settlement schemes for inter-provider connections (peerings) become a key issue [13]. For example, the costs of operation of the Internet for Swiss universities and research has to be recovered almost entirely from its users. This requires clearly defined methods for charging and accounting of Internet services as well as sound cost-sharing models and appropriate pricing schemes. Improvements on the existing model of cost sharing are required to achieve the level of detail for all involved cost factors of a given IP network. This includes various refinements to optimize resource usage in the Internet. In addition, various connected sites, organizations, and enterprises on the Internet have expressed a need for better cost attribution within their local administrative domain with respect to utilized IP services, resulting in proposals for charging systems being applicable to the Internet, e.g., [15], [34].

This paper is organized as follows. While Section 2 highlights pricing models in practice, Section 3 deals with a classification scheme for Internet pricing models in terms of scientific research. An overview of significant related scientific work can be found in Section 4. Research results and further investigations on dynamic market price implementations and simulations are discussed in Section 5. A concrete approach for the cost recovery of an ISP is discussed in Section 6. Finally, Section 7 summarizes the work and focuses on requirements on economically viable and technically efficient approaches for pricing models and cost recovery schemes as investigated in the framework of the Swiss National Science project CATI – Charging and Accounting Technology for the Internet [36].

2 Pricing Models in Practice

Only a few pricing schemes seem to be in wider use on the Internet today – especially from the ISP's point of view. Dial-up access to the Internet is often sold for a fixed monthly fee, including either unlimited use or a limited duration of the connection. Additional use of the connection is then billed at a per-hour fee. Volume-dependent charges have been used by some providers in the past, but seem to have lost popularity in recent years. The price of a fixed Internet connection is usually a monthly charge depending on the bandwidth of the connection. In many cases the customer has to pro-

vide the circuit to the Internet service provider's nearest point of presence. In most cases, this circuit will be a leased line, which adds another bandwidth-dependent recurring fee with an additional distance component.

For fixed connections, volume-based charging is more common, especially outside North America. There is usually still a fixed recurring component based on access capacity, but an additional cost per Megabyte (mostly with volume discounts) of data transferred over the connection. Such offers vary widely in the distribution between fixed and volume-dependent components, price/volume curves, and other parameters such as whether both directions of traffic are considered for the volume fee. In most cases, the bandwidth-based fixed part of the price includes some positive amount of volume that can be transmitted without charge.

Another variant that seems to have become popular recently is a “bursty” rate, where the Internet Service Provider periodically, *e.g.*, every hour, measures the volume of data transferred over the connection. For each charging interval, *e.g.*, every month, all samples are sorted by volume. A fixed percentage, *e.g.*, 5%, of the highest samples are discarded to eliminate unusual peaks, and the highest remaining sample is used to define the bandwidth at which the connection is charged.

It is quite rare to see further differentiation of usage-based charging, such as short and long distance, or peak and off-peak rates. Notable examples of those are volume-based charges on specific bottleneck links such as the trans-oceanic connections of New Zealand's and Britain's research networks [3], [27].

3 Pricing Model Classification

Since the variety of pricing models and schemes has grown in the past, a suitable comparison should be possible. Therefore, the proposed classification of pricing models must at least allow for (1) a clear service dependency in terms of the technical network type, (2) a suitable distinction of price components that need to be considered for pricing models, (3) a selection of parameters applicable to charging, and (4) the identification of pricing intentions.

3.1 Dimensions

In order to classify, characterize, and distinguish the different proposals for Internet pricing and charging mechanisms, three main dimensions have been identified: (1) the technical dimension, (2) the economic dimension, and (3) the research dimension. It has turned out additionally that each of them comprises two layers: one being higher in the sense of representing an abstract point of view, and the other one being lower in terms of dealing with more concrete issues. Overall, the resulting dimensions (cf. Figure 1) can be thought of as being more or less independent of each other, thus creating the domain within which any discussion on pricing and charging issues has to be placed.




Figure 1: Dimensions for Pricing Model Classification

3.1.1 Service Categories

Among existing proposals there is an apparent distinction between connection-oriented and connection less approaches. In the history of Internet pricing, *packet-based models* were the first to distinguish between several qualities of service, focussing mainly on single packets as being the entity to be charged. Quality-of-Service (QoS) levels are only designed in the form of relative best effort service classes. Adding guarantees to QoS classes has lead to the Differentiated Services Architecture (DiffServ) approach [1]. Here, the focus is on service classes characterized by guarantee parameters (abandoning the per-packet view or the per-flow view as defined within the Integrated Services Architecture [2]). DiffServ is commonly based on IP technology, *i.e.* connection less and hence at least questionable with respect to real-time applications. In contrast, the Integrated Services Architecture (IntServ) approach has changed to connection-like mechanisms for a sequence of packets of end-to-end associations – often called flow–, thus allowing high level QoS for these flows, as users now may be rejected during times of heavy traffic. But IntServ can no longer be run by IP only, since a signalling protocol like RSVP (Resource Reservation Protocol) is necessary. Note that neither IntServ nor DiffServ have taken over the Internet as known today, the Internet is still a single-service network offering best-effort type of services to all users. However, some overlay networks exist, which experiment on IntServ and DiffServ mechanisms.

Note that the terminology used above is only a specific one amongst a couple of other ones. *E.g.*, one may distinguish between *soft QoS* proposals (proposing different service categories for connection less traffic) and *hard QoS* proposals (focussing on real-time connection-oriented approaches only). Moreover, there are also a couple of ways to perform differentiation itself, *e.g.*, by priorities, price, capacity. For a recent survey see section 2.1 in [9].

3.1.2 Charging Parameters

This dimension considers the question of which technical parameters are or should be available for use in charging and pricing mechanisms, starting from priority flags and packet tagging over peak, nominal bit, or average flow rates as well as effective bandwidths to parameters like expected path and congestion cost or dynamic bid-prices per packet or resource unit, to name but a few. Further future investigations may encounter technical parameters such as delay, delay jitter, or error rates as main chargeable characteristics for, *e.g.*, real-time traffic as well. Certainly, the suitability of these parameters for a pricing scheme is limited by technical means of their monitoring or metering. Only those parameters which may be obtained efficiently may survive in the future.

3.1.3 Tariff Components

Common telecommunication pricing consists of three basic elements, *i.e.* access fees, setup fees, and usage fees [33]. Combining these three leads to a classification of pricing mechanisms, such as flat fee, usage-based, reservation-based, volume-based, service class-based, bandwidth-based. In the same way, there exist a couple of basic approaches for Internet pricing mechanisms to be distinguished, like flat rate, usage sensitive or volume-based pricing, packet or flow auctions, service or user profile-based mechanisms, edge pricing versus multilateral contracts.

3.1.4 Efficiency

Here, the basic function of pricing has to be investigated. Pricing may be used for reasons of network efficiency, *i.e.* maximizing the utilization of resources (*e.g.*, bandwidth, buffer space), or of economical efficiency, *i.e.* value to the user. Hence pricing maximizes either provider revenue (by efficient resource sharing and access control) or user satisfaction. Of course it should be noted that these types of efficiency are not necessarily orthogonal to each other (*i.e.* there may be models supporting network efficiency as well as economic efficiency), but nevertheless they represent two different aspects that have to be distinguished.

Incentive compatibility is a further aspect to be considered here, since it will allow for the definition of a significant relation between real users' resource demand and their willingness-to-pay.

3.1.5 Research Dimension

A last dimension describes, whether the research is conducted from a more theoretical or more practical and application-oriented point of view. Both types of research are able to deliver important input and insights especially for the development of an Internet architecture general enough for implementing different pricing models. For sure, this must remain a clear goal of Internet pricing work, since prices and their determination establish an important business aspect, which most certainly will not be similar comparing a number of ISPs.

3.2 Further Aspects and Requirements

Beyond such general dimensions, there are much more aspects and requirements to be considered while looking for viable pricing models, reaching from different types of applications (*e.g.*, burstiness issues) over technological and economical issues (*e.g.*, sender or receiver based payment, marginal cost, congestion/responsive pricing) to more practical ones (*e.g.*, like transparency, predictability, practicability, fairness, user acceptance, and user friendliness). Some of them will be investigated in detail later.

In addition, the work on pricing assumes that there are certain mechanisms and protocols in place to transfer charging information or transport money-related information. Questions on billing support in the traditional sense (*i.e.* collecting billing information, calculating the sum, printing an invoice, and sending it out to the user) need to be reconsidered due to performance reasons. Future systems may encounter a complete electronic billing system without any paper involved any more, *e.g.*, by using electronic fund transfers or credit cards. Or electronic payment systems may utilize micro-payment schemes for a full electronic handling on paying the prices and invoices for consumed communication services [33], [35].

4 Related Work on Pricing Models

Although the issue of pricing Internet services has gained its actuality only recently, there is already related work to be found in the scientific literature. First ideas during the early 90's usually preferred some simple sort of service priority models, until in 1994, [22] introduced the idea of using auction mechanisms. In 1995, [29] formulated an important model based on the IntServ approach. These two papers provided major stimuli to the research of the following years (*cf.* especially the collection of relevant papers in [23]), but it was not until 1997/98 that first steps towards designing "real world" charging mechanisms for IntServ took place, mainly based on the RSVP protocol [12], [14].

The following subsections contain an overview of important Internet pricing models which have been investigated over the last five years and have turned out to be of special importance from a practical and economic point of view. After presenting the idea of edge pricing, they deal with user and service profiles and concentrate on different types of auction mechanisms.

4.1 Edge Pricing

The fundamental idea of edge pricing [7], [30] is to charge the user only by the first network provider along a data path that might use also services from other providers. The charge to be paid includes expenses for all different providers handling the respective data. Thus, multilateral contracts are reduced towards a sequence of bilateral ones, the complexity is reduced enormously, and user transparency is provided. In the basic approach, the user defines the maximal total price she is willing to pay as a sender or a receiver of data, respectively, as well as an upper bound for the maximal number of hops. The charging information can be transmitted as part of a signalling protocol, *e.g.*, in the RSVP header [12], [14].

4.2 Profiles and Classes

[6] deals with the question of how to provide different QoS with high predictability while still running usual best effort. Instead of allocating capacity to users by explicit reservations, the “Expected Capacity” framework defines service profiles for each user and separates demand into those within profiles and those outside. Treating these two types of packets differently (*i.e.* favoring traffic that obeys the respective profile) allows the network to offer different levels of service with high predictability. In such a scheme, packets from a user behaving correctly are tagged “in”, whereas packets from a user exceeding her profile are tagged “out”. During congestion a suitable dropping scheme is used to preferentially drop “out” packets. Note that this approach prevents traffic being separated at routers into different flows or queues.

In contrast, [18] does not classify users, but services. Inside each service class every customer receives equal service, but higher service classes offer significantly better service than any lower service classes and, therefore, are charged higher prices. The Nominal Bit Rate (NBR) provides the underlying parameter for a monthly fee. Congestion is recognized by monitoring the load level of output buffers in the nodes; the system reacts by discarding some packets, preferably from flows with the actual bit rate to NBR ratio being high. Each packet carries drop preference and delay indication bits, based upon which the system decides about the discarding of packets.

The Paris Metro Pricing (PMP), another very interesting proposal [24], is based on subdividing the network into different logical subnetworks, each of them handling packets on a best-effort base, but charging different prices for them. This is an analogy to the price system used in the subway of Paris, and it is to be expected that a more expensive subnetwork will be frequented less often and is hence able to deliver high-quality service, but without giving formal guarantees for that.

4.3 Volume-based Schemes

As already noted in Section 2, volume-based charging applies prices to the amount of data transmitted. This concept has been applied commercially to X.25 networks as well as the different service classes of ATM traffic during the 1996 tariffs in Switzerland. A suitable metering component is required to monitor the amount of data transmitted. Most of these approaches use a system of price discounts based on several thresholds. For Internet traffic applying a volume-based scheme, the two examples of the traffic metering approach in New Zealand [3] and Great Britain [27] are already known. Note that current work in this area deals, *e.g.*, with stochastic network models that allow to derive so-called “price functions” describing the relationship between the current utilization of a resource and the price to be paid for using it [25].

4.4 Auction Mechanisms

The seminal work of [22] deals with the question of how an efficient pricing structure allows to manage congestion, encourage network growth, and guide resources to their most valuable use. As the marginal cost for transporting packets over the network is essentially zero as long as the network is not congested, usage-sensitive pricing schemes appear to be a good candidate for congestion control mechanisms, as they approach the allocation of scarce Internet resources in an economic context [17]. Note that the objective is not to raise profits, but to find a pricing mechanism yielding most efficient usage of existing resources. Current pricing schemes usually offer no incentives to flatten peaks or mechanisms for bandwidth allocation during congestion, whereas ideal prices should reflect resource costs that the user generates so that she can make informed decisions on resource utilization. Such costs include fixed costs for network infrastructure, costs of connecting to the network as well as sending extra packets, and finally the social costs of delaying other users’ packets during congestion periods.

A congestion pricing scheme (where packets are charged if and only if the network is congested) could be implemented by using a “smart market”, where the price for sending a packet varies on a very short time-scale, thus reflecting the current degree of network congestion [20]. Each packet header contains a bid field, and the packet is admitted, if the bid exceeds the current marginal cost of transportation. Note that the user does not pay the actual bid, but only the (lower) market-clearing price. In this kind of “second-bid” auction (Vickrey auction) the optimal strategy for the user is to bid her true evaluations – fooling the market results only in disadvantages. The mechanism guarantees only relative priority, no absolute QoS. Other critical issues include the question of how accounting should be done without yielding

too much of an overhead, how bursts will be handled, and how the user will react to maybe rapidly fluctuating bandwidth prices.

The model [20] allows guaranteeing multiple QoS (especially for inelastic traffic) by scheduling resources in advance. In order to maximize the sum of user utilities, a routing problem has to be solved by standard multi-commodity flow techniques. The notion of “effective bandwidth” allows to aggregate a broad range of source types in form of a one-dimensional bandwidth reservation [16].

De-coupling routing and usage optimization and solving the resulting linear problem and its dual, the latter allows interpreting the optimal solution in terms of spot prices for inserting (or extracting, respectively) traffic at a certain node. This allows to express the marginal system cost for traffic from node A to node B in terms of only two numbers: the nodal spot price for the source and for the sink of the flow; the user no longer needs to know about the optimal route. Note, however, that finding the optimal spot prices still requires solving the full central planning problem. Decentralization requires that users truthfully reveal their preferences so that a Pareto-efficient allocation can be calculated. The mechanism proposed is again a “smart market”, i.e. Generalized Vickrey Auctions (GVA), as presented in [22].

A related approach of using auctions as a proper method for decentralizing the decision-making in packet-switched multiservice integrated networks is presented in [19]. Here, Vickrey auctions are generalized yielding a mechanism that is designed to be stable, simple, efficient, and fair. While the original “smart market” proposal [22] uses one-dimensional bids (price per packet) and thus requires the central setting of the market clearing price based on explicitly assumed utility functions for the users, this approach requires the possibility of per-flow resource reservations. The resulting two-dimensional bids (price and quantity) allow to determine the clearing price directly from the bids only. Instead of dividing the resource into many small units and handling each of them as an indivisible object subject to a Vickrey auction (which yields a considerable loss of flexibility and scalability), here allocations are assumed for arbitrary shares of the total available resource quantity. Player preferences are given in the utility function describing the individual “value” of quantity/price vectors for the individual user. The Progressive Second Price (PSP) rule generalizes the idea of Vickrey auctions: you pay a price per unit which is calculated from all other players’ bids, where each of them is weighted by how much the allocation of that player is decreased by the existence of your bid. Hence, for each infinitesimal share of the resource, the player who is getting it, pays the maximum amount which the player who is denied would have been willing to pay for it. This rule can be shown to have a number of nice properties, ultimately leading to the existence of a fair and efficient Nash equilibrium.

Note that the application of auction mechanisms, as described above, to Internet scenarios that consist of multiple providers involved in establishing an end-to-end connection requires to deal with a couple of additional issues. These issues include at least:

- Synchronization problems arise from the fact that auctions take place locally per ISP at discrete moments which are independent of each other.
- Possible tear-downs of an end-to-end-connection may occur, because of loosing one local auction.
- The question need to be resolved of how to divide up a global user’s spending cap (i.e. a budget for a complete connection) into bids for local auctions on all individual parts of the end-to-end connection.

Dealing with these issues recently, has led to proposing two new auction schemes, i.e. Delta Auctions [12] and the CHIPS scheme [26]. These concepts will be introduced and investigated in Section 5.2 and 5.3.

Summarizing the approaches described, second price auctions appear to be a useful concept for determining actual market prices in case of network congestion which is one of the main prerequisites for a truly dynamic pricing scheme. However, it must be noted that according papers are kept rather theoretically and may need major adaptations before using their ideas for practical purposes. It has to be noted that the smart market model as well as auctions do show an often quoted drawback in terms of lacking price transparency and predictability, which results, *e.g.*, in problems for communication budget definitions. This is certainly true, however, as known and practically experienced, spot markets for oil, stock markets, and auctions of other goods are well established mechanisms in today’s trading environment. Therefore, an investigation and trial of user behavior and reaction on highly dynamic prices need to be performed to identify user acceptance for dynamic pricing schemes for Internet services. The INDEX project at Berkeley [10] conducted initial experiments in this domain of user behavior by focussing on static, usage-based pricing schemes.

5 Research Investigations and Results

For prototyping reasons and the availability of implementation environments, the Integrated Services Architecture (IntServ) of the Internet has been deployed and extended with required functionality for charging and accounting tasks. The Crossbow toolkit served as an experimental implementation platform for the required extensions [8]. As discussed in the following, results show that various pricing models can be integrated into existing Internet protocols, particularly the Resource Reservation Protocol (RSVP). Firstly, an explanation recalls basics of the extended IntServ implementation, secondly, new auction concepts for multiprovider scenarios are presented. Finally, shapes of market prices over time and the link utilization over time show results of auction simulations for dynamic pricing schemes.

5.1 A Charging Approach based on IntServ

Charging Internet services can be performed based on flows. Therefore, the Integrated Services (IntServ) Architecture offers the required basic protocol support. As shown in [12] and [32], the question of how to enhance a suitable signaling protocol, *e.g.*, the Resource Reservation Protocol (RSVP) with information for payment and prices can be answered. A set of new RSVP objects has been defined to carry required charging and pricing information between participating users and the network. In addition, depending on the underlying economic model, in particular the pricing model applied by an ISP, these information may be interpreted in different ways.

The basic idea is to use PATH messages (path messages are directed from senders to receivers and contain a sender-offered set of QoS specifications for a flow) and RESV messages (reservation messages are directed from receivers to senders and contain the requested QoS specification for one flow) of RSVP in order to transmit pricing information. Extended PATH messages carry a field with price information which is initially set to zero at the senders side. At each “hop” of an outgoing link, the current market price for the requested QoS is added to the price field, performed by an ISP on the selected path. The adding of the ISP’s portion of an end-to-end service price may be determined by an ISP’s local pricing scheme, valid for a dedicated section of the end-to-end connection only. In addition, the “hop” may be represented by a single router only, *e.g.*, an edge router of the entered ISP domain. Hence, if the PATH message has arrived at the receiver, crossing multiple ISPs, a current view of the market situation for this service is delivered to the receiver. Albeit, the final price may still vary slightly in case of dynamic pricing schemes, due to changes on the market situation. The RESV message is sent back and identifies the calculated price to the sender. After this round-trip time the reservation and data transfer phase may start as usually performed with RSVP. The extended RSVP messages may also be used for on-line payment information, *i.e.* the PATH messages may contain sender provided payments, whereas the RESV messages may carry receiver provided payments.

An initial prototypical implementation based on a first version of RSVP, called SSP (State Setup Protocol) for charging reserved flows has been performed.¹ The implementation is based on the flexible charging and accounting architecture as discussed in [11]. For experimentation purposes a static and a dynamic pricing scheme have been designed and implemented with respective charging functionality. This allowed for two sets of different measurements being performed using two different pricing models, *i.e.* Dynamic Volume Pricing and Delta Auctions, for comparisons. Delta auctions avoid the generation of huge signaling message bursts (as traditional second price auctions would do) due to distributing the determination of an auction result at a single ISP over a fixed and small auction time [26]. Furthermore, users are informed early of rejected bids for a flow. Therefore, the following section introduces this new auction concept in more detail and presents its enhancement for multiproviders called CHiPS, before a number of measurement and simulation results will be presented afterwards.

5.2 New Auction Concepts for Multiprovider Internet Scenarios

As already mentioned in Section 4.4, applying auctions to scenarios where a connection has to cross more than one ISP, leads to a number of important questions, especially concerning the synchronization of individual auctions which have to be won completely in order to establish the end-to-end connection, the influence of local market instabilities with respect to the global connection, and the question of how to divide the global user spending cap into bids for the sequence of individual auctions. This section presents some rather new concepts which aim at solving such difficulties.

Assume a connection to be established comprises n local auctions at different ISPs i with respective auction periods (*i.e.* reservation intervals) of P_i . Then, missing synchronization between these auctions yields a mean delay of $(n/2) \sum P_i$ until the establishment of a connection. The basic idea of Delta auctions aims at making the auction procedure continuous by immediately processing arriving requests and rejecting bids that will not have a chance as soon as possible. In this way, it is possible to smoothen signalling traffic and to prevent users from idle waiting for auction results that are of no use to them. Further details of this concept are to be found in [12].

As soon as an end-to-end connection has been established successfully, usual resource reservation protocols require reservations to be refreshed and, hence, the auctions to be repeated periodically. In this case, rapidly changing market conditions may cause local auctions to be lost, thus, yielding the reservation for a may be small part of the connection to be lost. Consequently, the complete connection is to be torn down. CHiPS (the “Connection-Holder-is-Preferred Scheme”) deals with this undesirable situation by giving holders of a connection the chance to increase their bid for the lost auction a posteriori such that the bid now outnumbers the current market price once again. This is possible, because the user’s spending cap in effect does not limit his total bid sum, but the total sum she may spend (which under second price auctions usually is lower than the bids). Moreover, it is made sure that no user has to pay more than she previously has agreed to and that the second price character of the auction is essentially preserved.

Finally, a new bid structure has been introduced within CHiPS, in order to take into account local auctions especially endangered, because of market conditions which may vary more than usual. Moreover, this new bid structure, called “AMF Scheme”, also provides an elegant way to signal that a user is willing to increase her bid for a local auction *ex post* after having lost this auction.

1. At the TIK laboratory there is a JAVA-based, unicast RSVP implementation in progress, including charging extensions.

To this end, assume that the amount a user is bidding for a specific local auction, consists of three parameters: the current market price at the auction place m , a global bid factor f calculated by the user and essentially representing the ratio of the user spending cap and the global sum of current market prices, and a parameter a held locally at each auction place and initially set to zero. Hence, the complete bid for a local auction i looks like

$$b_i = a_i + m_i \cdot f, \quad (1)$$

thus, hinting on the origin of the name ‘‘AMF’’. Now imagine a local auction gets lost for a user already holding a connection. According to CHiPS, this user needs to be informed that she will have to increase her bid for that specific auction ex post in order to maintain her connection. This may take place in two different ways: either a RSVP message is sent to the user stating that she will have to increase her bid by an amount of d , say, whereupon the user sends a message back to that auction confirming or refusing this increase. The second possibility to signal the auction loss and the increase, respectively, works as follows: the auction sends an RSVP message in both directions, i.e. to both end-points of the connection, decreasing the local parameter a at all intermediate auction places by d . In case the user agrees with increasing her bid, she sends an RSVP message to the receiver which now increases all a -parameters on its way (now including the lost auction) again by d . The effect of this procedure is that the a -parameters of all auctions remain unchanged except for the one auction that has previously been lost. For the latter one, the bid now is increased by d as required, moreover, in future each bid at that auction will automatically be slightly higher than the average (represented by the bid factor f), thus taking into account that the market fluctuation at this places has turned to be likely higher than elsewhere). For further details of this scheme refer to [26]. Moreover, the following Section 5.3 presents some simulation results obtained by implementing this scheme as well as implementation measurement results of the IntServ-based charging approach with RSVP extensions.

5.3 Evaluation and Simulation of Dynamic Pricing Schemes

The overhead of the charging and accounting extensions due to the inclusion of charging functionality were measured in a prototypical implementation with relation to the requested resources. In this case, required bandwidth, necessary memory for accounting storage, and needed processing within SSP routers were used as main measures. All measurements have been performed on an Integrated Services testbed (Crossbow) of routers and hosts [8], which included standard PCs with 233 MHz AMD K6 processors running NetBSD 1.21 and applied an IP telephony application and its data flow for reservation. With 5 s long reservation periods, 80 Byte telephony packets, and a G.711 type of traffic, a sort of worst case situation has been investigated.

An overview and comparison is indicated in Figure 2. In this case the memory needed to store flow states extended by charging and pricing information at each router encompasses 194 Byte in the SSP daemon in user space, where the flow state requires 120 Byte, pricing 14 Byte, the account 12 Byte, and basic authentication data of 28 Byte is also included in the form of an integrity object with an MD5 checksum, sequence number, previous hop, and time stamp. Therefore, the charging relevant data uses 28% of the space for the total SSP flow state information.




Figure 2: Charging Overheads for IntServ

Processing at SSP routers to perform price queries, the delta auction or a volume-based scheme, and authentication requires for a long living flow on the average 322 μ s per packet which encounters for 1.8% for the sample scenario of IP Telephony flows.

A certain protocol overhead is needed to transport charging relevant information. Again, as an example the same IP telephony stream has been used, with fine grained reservation or refresh periods of 5 seconds and a constant bit rate of 64 kbit/s (40,000 Byte telephony data in 500 data packets for this 5 second period). However, charging information for this stream uses only 0.8% of the flow’s total bandwidth, including the signalling messages. For a single reservation period a message round-trip time for querying the price and refreshing the reservation is needed.

Depending on the user data bandwidth requirements per flow, these protocol overheads vary. As shown in Figure 3, the overhead for 5 s refresh periods, and equally achievable a 5 s billing granularity in turn, remains even for small flows, such as a 2.4 kbit/s stream, below 1%.

Comparing the two different pricing models applied in the implementation (Delta Auction and Dynamic Volume Pricing), Figure 4 shows processing times for a single reservation request including the message processing of the charging information. These interface-to-interface delays on a router connected to two 10 Mbit/s Ethernet subnetworks show delays for a reservation message of 150 μ s and 144 μ s for the Delta Auction and the Dynamic Volume Pricing model,

respectively. While the latter does add a minimal processing delay (<0.7% of the total delay), the more complex Delta Auction uses less than 6.7% of the overall delay. Note, the IP protocol stack processing delay including the protocol overhead, the required socket processing, and the use of timers, remains similar at 126 μ s and 125 μ s for both pricing schemes.




Figure 3: Charging Overheads Depending on the Refresh Period Length

These performance evaluations show for a basic SSP implementation, which is functionally similar to an RSVP implementation, and its extensions with charging and accounting functionality that the IntServ architecture can be extended efficiently for flow-based charging purposes. However, at this point the problems of IntServ’s scalability issues have not been dealt with. Future work is directed towards the integration of charging and accounting functionality into the DiffServ architecture and to be worked at signaling protocols in the same environment.

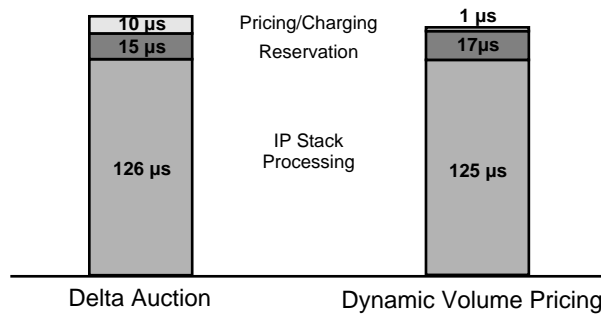


Figure 4: Processing Times for Pricing Schemes

Having demonstrated that the use of a dynamic pricing scheme does not imply intolerable processing and protocol overheads, the remainder of this section presents simulation results concerning the new auction concepts introduced in Section 5.2. These concepts have been evaluated in further steps, particularly, the development of the dynamic market price over time, based on the link’s utilization. Therefore, a simulation has been set up and a scenario with non-congested and congested links has been simulated. These investigations are necessary to extend performance results of single end-system and router behaviors to a networked scenario, where the course of dynamic market prices can be analyzed. The simulation has been performed using an extended version of FlowSim being developed at the ETH Zürich [28], [31]. FlowSim is Java-based and has been developed as a small, understandable, and fast simulation tool based on the concept of “flows” as the smallest unit to be simulated and thus being in contrast to well-known packet-based tools like ns-2 [38].

Figure 5 shows the scenario used in the simulation which is able to comprise a multitude of interesting aspects. At this point, the focus is directed towards the bottleneck link between nodes 4 and 6. The bandwidth of regular links is assumed to be 155 Mbit/s, whereas in the bottleneck case the bandwidth has been reduced in different steps down to 2 Mbit/s.




Figure 5: Basic Simulation Scenario

This scenario has been used to investigate a couple of important constellations. In a first step, it has been assumed that there is a user (*e.g.*, node 0) whose spending cap is significantly higher than the budget of the other ones (*e.g.*, nodes 1-3). For this case, Figure 6 left shows the utilization of user 0 (in kbit/s) and the market price on the bottleneck link. It is shown, that on one hand user 0 indeed rules out other users (occupying nearly always the complete available bandwidth) while on the other hand Figure 6 right demonstrates that the resulting market price still may vary by large amounts.




Figure 6: Utilization and Market Price on the Bottleneck Link

In a next step the influence of variable spending caps has been investigated. Figure 7 shows the resulting market dynamics. Here, it has been assumed that the spending cap lies either between 1000 and 1010 monetary units (upper figure), between 1000 and 1050 (middle figure), or between 1000 and 1100 (lower figure), i.e. varies by 1%, 5%, or 10%, respectively. The individual bid has been chosen randomly within these boundaries. From these figures it can be derived that the market price tends to be quite often near the upper limit of the spending caps, but may also assume amounts throughout the whole range of possible market prices.




Figure 7: Market Price on the Bottleneck Link While Spending Caps Rise

As a last example, a first result is presented concerning especially the CHiPS mechanism as described in Section 5.2. In this case, it is of special interest, if CHiPS is able to preserve the incentive compatibility (the main characteristic for Second Price Auctions) in the multiprovider case. To this end, different bidding strategies for a user have been simulated, depending on the true evaluation of a service. The easiest way to parameterize these strategies is by using a so-called “strategy factor” s that relates the true value of the resource to the chosen bid, e.g., $s=0.5$ corresponds to bidding only half of the value the resource has indeed for the user. Using this parameter, it is possible to determine the utility for the user (i.e. the difference between what she gets and what she pays) depending on the strategy factor s . Figure 8 presents a typical result indicating that bidding exactly the true value of the services (which is equivalent to a strategy factor 1.0 on the x-axis) yields an optimum in terms of the utility function (y-axis). Such results allow to conclude that the CHiPS approach indeed is incentive compatible.




Figure 8: CHiPS' Incentive Compatibility

6 Cost Recovery within SWITCH

SWITCH is the Internet service provider for research and higher education in Switzerland and is organized as a non-profit foundation with little centralized funding. In yearly budget discussions, representatives of foundation members define a charging scheme for the next year. These charges are set to recover all operating costs for the network.

In recent years, the tariff for leased line connections consisted of an access charge based on the capacity of the connection, and a charge per gigabyte of data transferred from SWITCH towards the connected organization. The fixed charge and volume-based charge were set to account for about one third and two thirds of the total, respectively.

The reasons for only counting traffic towards customers are as follows: The underlying transmission links of the network are always sold symmetrically, but more traffic flows towards the universities than from them. Thus, there is usually no congestion in the outbound direction. Secondly, the pricing is set to encourage sharing of information.

The large part of usage-based charges has caused a high level of awareness of costs for the network usage within the organizations connected to SWITCH. Several creative methods have been devised by universities to reduce the amount of traffic and control the cost associated with their network connection. This includes reducing WWW bandwidth consumption by use of caching proxy servers (in some cases made mandatory through blocking of direct connections), or converting USENET News feeds from “push” to “pull” mode.

Some of these measures to reduce traffic volume have actually been counterproductive in the sense that they reduced traffic only on those lines that were not congested anyway, so they did not actually reduce cost of operation. In some cases, efforts to save volume even made the service more expensive for everyone, because of the increased load on servers: In the traditional “push” feed model, articles can be transferred to downstream servers as they are received. In the “pull” model that was adopted by many connected organizations to reduce volume-based charges, downstream servers request articles asynchronously. This has adverse effects on locality of reference in the upstream server process, thus, increasing its resource requirements.

These concerns have been addressed in the charging scheme for the year 2000 by the following modifications: The volume-dependent share has been reduced to one third of the total fee. Volume charges only apply to traffic over the transatlantic US line, so internal traffic and traffic over cost-neutral peering connections is no longer penalized. Finally, an off-peak discount of 75% applies during nights and over the weekend. The hope is that this more differentiated tariff leads to more efficient network usage which outweighs the cost of a more complex accounting scheme.

Another tariff scheme will be introduced with the addition of some twenty new sites to the academic network. The new organizations will pay flat rates depending on two bandwidth quantities: the total capacity of their access circuit into the network and a “committed” rate directed to the US which can be chosen at any rate up to the total access capacity. The fee associated with this committed US rate is set to correspond to the cost of upgrading SWITCH’s US connection by that amount. Rate guarantees are provided through provisioning and a DiffServ-like configuration on routers terminating the US-to-Switzerland link. Traffic for a given customer is metered against the committed rate profile, while conforming packets being marked with higher than normal priority and excess packets with lower than normal priority. Weighted Random Early Discard is used to selectively discard lower-priority traffic in times of congestion.

SWITCH participants have often expressed their intention to charge individual users (or organizational units such as departments or institutes) for the volume of network traffic generated. So far, the technical and administrative complexity involved with this has prevented them from doing so. A few universities provide tools through which individual users can inform themselves about their amount of network usage [37]. This has been found very useful and could obviate the need of actually performing accounting and billing within such organizations.

6.1 Cost Structure of SWITCH’s Services

In many respects, SWITCH is not exactly representative for Internet Service Providers: the essentially closed user group, organizational structure as a non-profit foundation, and emphasis on communication needs of the academic community make it different from most ISPs which operate on a commercial basis and offer services to a wide range of individual or corporate users in a very competitive market. However, the cost structure of operating SWITCH’s network service is quite similar to those of other large providers with mainly leased-line customers. SWITCH operates a national backbone (cf. Figure 9) to transport data between connected sites. This national backbone is built from routers operated by SWITCH, connected with leased lines or ATM (Asynchronous Transfer Mode) connections. For both types of connections, SWITCH pays monthly volume-independent fees to carriers.

For traffic between SWITCH-connected sites and other networks, there are two different cases: Some providers’ networks are reachable at low cost, because a “peering agreement” permits mutual exchange of traffic at no charge. In those cases, router ports and lines connecting the two networks are the only costs involved. Where peering agreements cannot be established, SWITCH has to pay other networks for “transit” to other parts of the Internet. Currently, SWITCH uses transit service from a commercial ISP in the U.S.A. and leases its own line (actually an ATM virtual circuit) to New York. This transit subscription and leased line represents a very significant fraction of the cost of operating SWITCH.

Because international circuits are much more expensive than domestic ones, the costs for international connectivity exceed the costs for the national backbone. Since international transit connections are used to ensure connectivity to

those parts of the Internet that cannot be reached by a peering agreement, there is a strong incentive to peer with other networks to which there is significant traffic.

6.2 Peering Agreements

This explains why Internet Service Providers agree to exchange mutual traffic without settlements. However, such agreements are hindered by the fact that the savings are usually not perfectly balanced, i.e. one of the potential peers would benefit from the peering more than the other. An ISP who believes that the potential peer would reap larger benefits may see the peering as counter-productive, because it would make the other ISP more competitive at the first one's expense. In cases where a difference in benefits is perceived, peering parties usually negotiate some kind of compensation, such as one party paying for the entire line between the two networks, or the provision of free transit to a third network by one of the potential peers. Where the benefits are seen as widely differing, the provider with only small perceived benefits will likely not go through the effort of setting up the peering.

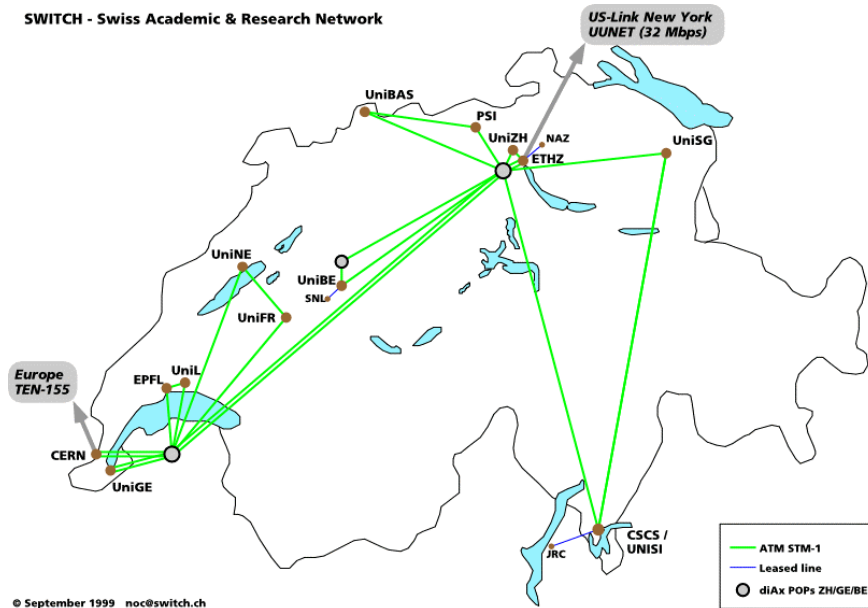


Figure 9: SWITCH's National Backbone

Assessing the mutual benefits of existing peerings is often done, when upgrades have to be negotiated. The metric that is used most often is the amount of traffic flowing in both directions. Traffic sent from ISP A to ISP B is generally thought of as beneficial to ISP B, because presumably it represents data that has been requested by ISP B's customers. This crude measure still corresponds to the prevalent ways the Internet is used today. However, it may be challenged by the emergence of unsolicited information such as bulk e-mail or paid advertisements.

While this metric is useful enough to serve in arguments on establishing or upgrading settlement-free peering agreements, it would probably not be sufficient as a basis for settlements between peering ISPs. In fact, peering with settlements does not seem to be popular at all. The only scheme used in practice is, that one (smaller) ISP pays another (larger) one to connect to his network, but without the right to transit to other networks. The authors are not aware of cases, where settlements between peers are set in both directions and based on a common metric.

7 Summary and Remaining Issues

This paper provides a review of a large variety of possible pricing schemes that could be used in the context of Internet pricing and cost models. In addition, basic research results in terms of a usage-sensitive charging approach for multimedia flows have been sketched to demonstrate by a prototypical implementation and simulations that flow-based charging is possible in the Internet, including an incentive compatible pricing scheme approach. Moreover, major relevant issues concerning cost structures and cost recovery have also been enlightened from an Internet Service Provider's point of view. This is especially important for judging the practical relevance of research work in a real business environment with daily customers and the handling of their requirements.

As most of the work on pricing schemes so far is kept rather theoretically, the next step is to validate their practical applicability. Having future general Internet developments in mind, in the authors opinion an ideal implementation of a pricing model requires the essential integration of the following concepts:

- Sufficient support of temporal pricing aspects;
- Using some sort of second-price auction approach for dynamic pricing, but refining it in a way that the auctions and, therefore, the final price become more transparent to the user;

- Simplifying tariffs by offering a suitably defined traffic classification;
- Extending pricing models and the corresponding technical protocol-based signalling to the case of multiprovider networks;
- Keeping pricing and charging issues as far as possible at the border of the ISP's networks;
- Investigating a suitable granularity of dynamically charged services in terms of technical service parameters as well as timing informations (e.g., duration).

Creating a platform that allows for testing and using different pricing models in a practical multiprovider environment – integrating both the IntServ and DiffServ worlds – is one of the major goals of the Swiss National Science Foundation's project CATI [5], [36]. The prototype currently developed focuses on an implemented charging and accounting framework to be used, e.g., for tariffing IP telephony, as well as for DiffServ-based Virtual Private Network configurations. The cooperation with the Swiss ISP SWITCH thereby guarantees that the project results will fit closely the features required by practical users of an Internet charging and accounting tool.

Acknowledgments

This work has been performed in the framework of the project Charging and Accounting Technology for the Internet – CATI (CAPIV 5003-054559/1 and MEDeB 5003-054560/1) which has been funded by the Swiss National Science Foundation, SNF, Bern, Switzerland. The authors like to acknowledge B. Plattner and T. Braun as main project contractors, valuable contributions of their colleagues G. Fankhauser and W. Huber, and the student's work of B. Spielmann and F. Floessl. Finally, thanks are addressed to the anonymous referees for a very constructive list of suggestions which helped to improve the quality of this paper.

References

- [1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss: *An Architecture for Differentiated Services*; Request for Comments RFC 2475, Internet Engineering Task Force, December 1998.
- [2] R. Braden, D. D. Clark, S. Shenker: *Integrated Services in the Internet Architecture: An Overview*; Request for Comments RFC 1633, Internet Engineering Task Force, June 1994.
- [3] N. Brownlee: *Internet Pricing in Practice*; In: [23] pp 77-90.
- [4] G. Carle, M. Smirnov, T. Zseby: *Charging and Accounting Architectures for IP Multicast Integrated Services over ATM*; 4th International Symposium on Interworking (Interworking'98), Ottawa, Canada, July 1998.
- [5] *CATI – Charging and Accounting Technology for the Internet*; Up-to-date information available at the URL: <http://www.tik.ee.ethz.ch/~cati>, 1999.
- [6] D. Clark, W. Fang: *Explicit Allocation of Best-Effort Packet Delivery Service*; IEEE/ACM Transactions on Networking, Vol. 6, No. 4, August 1998.
- [7] 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.
- [8] D. Decasper, M. Waldvogel, Z. Dittia, H. Adishesu, G. Parulkar, B. Plattner: *Crossbow – A Toolkit for Integrated Services over Cell-Switched IPv6*; IEEE ATM'97 Workshop, Lisboa, Portugal, June 1997.
- [9] C. Dovrolis, D. Stiliadis, P. Ramanathan: *Proportional Differentiated Services: Delay Differentiation and Packet Scheduling*; ACM Computer Communication Review (SIGCOMM'99), Vol. 29, No. 4, October 1999, pp 109–120.
- [10] R. J. Edell, P. P. Varaiya: *Providing Internet Access: What we can learn from the INDEX trial*. Keynote talk at Infocom'99, New York, March 1999. INDEX Project Report No. 99-010W available at the URL: <http://www.index.berkeley.edu>.
- [11] G. Fankhauser, B. Stiller, B. Plattner: *Arrow: A Flexible Architecture for an Accounting and Charging Infrastructure in the Next Generation Internet*; Netnomics, Vol. 1, No. 2, 1999, pp 201–223.
- [12] G. Fankhauser, B. Stiller, C. Vögtli, B. Plattner: *Reservation-based Charging in an Integrated Services Network*; 4th INFORMS Telecommunications Conference, Boca Raton, Florida, U.S.A., March 1998.
- [13] G. Huston: *Interconnection, Peering, and Settlements*; The Internet Summit (INET'99), San Jose, California, U.S.A., June 22–25, 1999, E-commerce and E-business Track 1, Session on the Politics and Economics of Peering and Interconnection.
- [14] 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.
- [15] N. Kausar, B. Briscoe, J. Crowcroft: *A Charging Model for Sessions on the Internet*; 5th European Conference on Multimedia Applications, Services, and Techniques (ECMAST'99); Madrid, Spain, May 26–28, 1999, LNCS Vol. 1629, Springer, Berlin, Germany, pp 246–261.
- [16] F. Kelly: *Tariffs and Effective Bandwidths in Multiservice Networks*; In “The Fundamental Role of Teletraffic in the Evolution of Telecommunication Networks”, J. Labetoulle and J. W. Roberts (eds.), Elsevier Publishers, Amsterdam, The Netherlands, 1994, pp 401 – 410.

- [17] F. Kelly, A. Maulloo, D. Tan: *Rate Control in Communication Networks: Shadow Prices, Proportional Fairness and Stability*; Journal of the Operational Research Society, Vol. 49, 1998, pp 237–252.
- [18] K. Kilkki et al.: *Internet Charging Reconsidered*; 5th Annual Network and Interop Engineers Conference (Interop'98), Las Vegas, Nevada, U.S.A., May 1998.
- [19] A. Lazar, N. Semret: *Auctions for Network Resource Sharing*; CTR Technical Report, Columbia University New York, U.S.A., February 1997.
- [20] J. MacKie-Mason: *A Smart Market for Resource Reservation in a Multiple Quality of Service Information Network*; Technical Report, University of Michigan, U.S.A., September 1997.
- [21] J. MacKie-Mason, J. Murphy, L. Murphy: *The Role of Responsive Pricing the Internet*; Technical Report, University of Michigan, U.S.A., June 1995. Also in: [23].
- [22] J. MacKie-Mason, H. Varian: *Pricing the Internet*; In Public Access to the Internet, B. Kahn, J. Keller (eds.), Prentice Hall, Englewood Cliffs, New Jersey, U.S.A., 1995.
- [23] L. McKnight, J. Bailey (eds.): *Internet Economics*; MIT Press, Cambridge, Massachusetts, U.S.A., 1997.
- [24] A. M. Odlyzko: *Paris Metro Pricing for the Internet.*; ACM Conference on Electronic Commerce (EC'99), Denver, Colorado, U.S.A., November 3–5, 1999.
- [25] P. Reichl: *Kelly's Bound, RUAA and the Pricing of Multiclass Traffic in Loss Networks*; UK Performance Engineering Workshop, Bristol, U.K., July 1999.
- [26] P. Reichl, G. Fankhauser, B. Stiller: *Auction Models for Multiprovider Internet Connections*; Messung, Modellierung und Bewertung von Rechensystemen (MMB'99), Trier, Germany, September 21–22, 1999.
- [27] R. Rogerson: *Usage-related Charges for the JANET Network*; JISC Circular 3/98, March 1998, available at the URL: http://www.jisc.ac.uk/pub98/c3_98.html.
- [28] D. Schweikert: *QoS Routing and Pricing in Large Scale Internetworks*; Diploma Thesis, TIK, ETH Zürich, Switzerland, March 1999.
- [29] S. Shenker: *Fundamental Design Issues for the Future Internet*; IEEE Journal on Selected Areas in Communications, Vol. 13, No. 7, September 1995, pp 1176–1188.
- [30] 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.
- [31] B. Spielmann: *Design and Evaluation of Efficient Pricing Schemes for Integrated Internet Services*; Diploma Thesis, TIK, ETH Zurich, Switzerland, July 1999.
- [32] B. Stiller, G. Fankhauser, B. Plattner: *Charging Multimedia Flows in an Integrated Services Network*; 8th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'98), July 8–10, 1998, Cambridge, England, U.K., pp 189–192.
- [33] B. Stiller, G. Fankhauser, B. Plattner, N. Weiler: *Pre-study on Customer Care, Accounting, Charging, and Pricing*; Pre-study for the Swiss National Science Foundation, Bern, Switzerland, February 1998, available at the URL: <ftp://ftp.tik.ee.ethz.ch/pub/people/stiller/pre-study/all.ps.gz>.
- [34] B. Stiller, G. Fankhauser, B. Plattner, N. Weiler: *Charging and Accounting for Integrated Internet Services – State of the Art, Problems, and Trends*; The Internet Summit (INET'98), Geneva, Switzerland, July 21 – 24, 1998, Commerce and Finance Track 3, Session on Commerce Issues for ISPs.
- [35] B. Stiller, G. Fankhauser, G. Joller, P. Reichl, N. Weiler: *Open Charging and QOS Interfaces for IP Telephony*; The Internet Summit (INET'99), San Jose, California, U.S.A., June 22–25, 1999, Technology Track 4, Session on IP Audio.
- [36] B. Stiller, T. Braun, M. Günter, B. Plattner: *The CATI Project – Charging and Accounting Technology for the Internet*; 5th European Conference on Multimedia Applications, Services, and Techniques (ECMAST'99); Madrid, Spain, May 26–28, 1999, LNCS Vol. 1629, Springer, Berlin, Germany, pp 281–296.
- [37] R. Timsit, J. Virchaux: *MATHE (Mesure et Analyse du Trafic Hors EPFL)*; Flash Informatique 4/1998, Lausanne, Switzerland, April 1998, available at the URL: <http://www.epfl.ch/SIC/SA/publications/FI98/fi-4-98/4-98-page5.html>.
- [38] UCB/LBNL/VINT: *Network Simulator ns (version 2)*; <http://www-mash.cs.berkeley.edu/ns>.